DOCUMENT RESUME

ED 482 045 IR 058 778

AUTHOR	Bearman, David, Ed.; Trant, Jennifer, Ed.
TITLE	Museums and the Web 2002: Selected Papers from an International Conference (6th, Boston, Massachusetts, April 17-20, 2002).
INSTITUTION	Archives and Museum Informatics, Pittsburgh, PA.
ISBN	ISBN-1-885626-25-1
PUB DATE	2002-04-00
NOTE	252p.; Accompanying CD-ROM is not available from ERIC. For selected individual papers, see IR 058 779-800.
AVAILABLE FROM	Archives & Museum Informatics, 2008 Murray Ave., Suite D, Pittsburgh, PA 15217. Tel: 412-422-8530; Fax: 412-422-8594; e-mail: info@archimuse.com; Web site: http://www.archimuse.com/. For full text: http://www.archimuse.com/mw2002/ .
PUB TYPE	Collected Works - Proceedings (021) -- Non-Print Media (100)
EDRS PRICE	EDRS Price MF01/PC11 Plus Postage.
DESCRIPTORS	*Information Technology; Multimedia Materials; *Museums; Technology Uses in Education; User Needs (Information); Users (Information); *Web Sites; *World Wide Web
IDENTIFIERS	Electronic Resources; *Virtual Museums; Web Site Design

ABSTRACT

This proceedings contains the following selected papers from the Museums and the Web 2002 international conference: "The Electronic Guidebook: Using Portable Devices and a Wireless Web-Based Network To Extend the Museum Experience" (Robert Semper, Mirjana Spasojevic); "Eavesdropping on Electronic Guidebooks: Observing Learning Resources in Shared Listening Environments" (Allison Woodruff, Paul M. Aoki, Rebecca E. Grinter, Amy Hurst, Margaret H. Szymanski, James D. Thornton); "Can You See Me? Exploring Co-Visiting Between Physical and Virtual Visitors" (Areti Galani, Matthew Chalmers); "The Museum Wearable: Real-Time Sensor-Driven Understanding of Visitors' Interests for Personalized Visually-Augmented Museum Experiences" (Flavia Sparacino); "Now That We've Found the 'Hidden Web,' What Can We Do With It? The Illinois Open Archives Initiative Metadata Harvesting Experience" (Timothy W. Cole, Joanne Kaczmarek, Paul F. Marty, Christopher J. Prom, Beth Sandore, Sarah Shreeves); "Combining the CIDOC CRM and MPEG-7 To Describe Multimedia in Museums" (Jane Hunter); "Today's Authoring Tools for Tomorrow's Semantic Web" (Andy Dingley, Paul Shabajee); "The Virtual Ramp to the Equivalent Experience in the Virtual Museum: Accessibility to Museums on the Web" (Liddy Nevile, Charles McCathieNevile); "Adding Value to Large Multimedia Collections through Annotation Technologies and Tools: Serving Communities of Interest" (Paul Shabajee, Libby Miller); "Content Management for a Content-Rich Website" (Nik Honeysett); "Here and There: Managing Multiply-Purposed Digital Assets on the Duyfken Web Site" (Marjolein Towler, Valerie Hobbs, Diarmuid Pigott); "Pyramid Power: A Train-the-Trainer Model To Increase Teacher Usage of the ArtsConnectEd On-Line Resource" (Scott Sayre, Kris Wetterlund); "Digital Primary Source Materials in the Classroom" (Nuala Bennett, Brenda Trofanenko); "Statistics, Structures & Satisfied Customers: Using Web Log Data To Improve Site Performance" (Darren Peacock); "How Do You Like To Learn? Comparing User Preferences and Visit Length of Educational Web Sites" (David T. Schaller, Steven Allison-Bunnell, Minda Borun, Margaret B.

Reproductions supplied by EDRS are the best that can be made
from the original document.

Chambers); "Evaluating The Features of Museum Websites: (The Bologna Report)" (Nicoletta Di Blas, Carolina Orsini, Maria Pia Guermandi, Paolo Panlini); "Towards Tangible Virtualities: Tangialities" (Slavko Milekic); "Making It Realtime: Exploring the Use of Optimized Realtime Environments for Historical Simulation and Education" (Chris Calef, Turlif Vilbrandt, Carl Vilbrandt, Janet Goodwin, James Goodwin); "Networked Multi-Sensory Experiences: Beyond Browsers on the Web and in the Museum" (Fabian Wagmister, Jeff Burke); "Systematically Speaking: How Do Natural History Museum Web Sites Represent Science?" (Roy Hawkey); "Hacking Culture" (Pia Vigh); and "Storytelling and the Web in South African Museums" (Katherine J. Goodnow, Yngvar Natland). Author biographies are included. An accompanying CD-ROM includes: a list of all the speakers at the conference and links to their abstracts, biographies, and papers (where available); an overview of the Museums and the Web 2002 conference program and links to abstracts and paper biographies; and the results of the Best of the Web 2002 conference (requires Internet access). Most papers contain references. (MES)

Reproductions supplied by EDRS are the best that can be made
from the original document.

IR058778

Museums and the Web 2002

Selected Papers from an International Conference

PERMISSION TO REPRODUCE AND
DISSEMINATE THIS MATERIAL HAS
BEEN GRANTED BY

D. Bearman

TO THE EDUCATIONAL RESOURCES
INFORMATION CENTER (ERIC)

1

U.S. DEPARTMENT OF EDUCATION
Office of Educational Research and Improvement
EDUCATIONAL RESOURCES INFORMATION
CENTER (ERIC)

☒ This document has been reproduced as received from the person or organization originating it.

☐ Minor changes have been made to improve reproduction quality.

- Points of view or opinions stated in this document do not necessarily represent official OERI position or policy.

edited by
David Bearman and
Jennifer Trant

BEST CO

www.archimuse.com/
mw2002/

Museums
and the
Web 2002

Archives & Museum Informatics

Consulting, Publishing and Training for Cultural Heritage Professionals

***Informatics**: The interdisciplinary study of information content, representation, technology and applications and the methods and strategies by which information is used in organizations, networks, cultures and societies.*

Archives & Museum Informatics organizes conferences, workshops and seminars, publishes journals and monographs, and consults for cultural heritage organizations worldwide. For over 10 years, our educational goal has been to provide specialists in archives, museum and cultural and heritage informatics with timely and challenging opportunities for professional exchange and training. Our consulting services emphasize inter-institutional collaboration,strategic planning and standards-based solutions.

To be kept informed of coming events, conferences and new publications contact us at:

Archives & Museum Informatics
2008 Murray Avenue, Suite D
Pittsburgh, Pennsylvania
15217 USA
phone: +1 412 422 8530
fax: +1 412 422 8594
email: info@archimuse.com
http://www.archimuse.com

Ask to be put on our mailing list.

www.archimuse.com/
mw2002/

Museums and the Web 2002

Selected Papers from an International Conference

edited by
David Bearman and
Jennifer Trant

Museums and the Web 2002
Selected papers from an international conference

edited by David Bearman and Jennifer Trant
Pittsburgh, PA: Archives & Museum Informatics
ISBN: 1-885626-25-1

Printed in the United States of America

Contents

Introduction

Cyberspace in Our Space
David Bearman and Jennifer Trant, Archives & Museum Informatics, USA 3

The Edge of the Web

The Electronic Guidebook: Using Portable Devices and a Wireless Web-based Network to Extend the Museum Experience
Robert Semper, Exploratorium, and Mirjana Spasojevic, Hewlett-Packard Labs, USA 11

Eavesdropping on Electronic Guidebooks: Observing Learning Resources in Shared Listening Environments
Allison Woodruff, Paul M. Aoki, Rebecca E. Grinter, Amy Hurst, Margaret H. Szymanski, and James D. Thornton, Palo Alto Research Center, USA 21

Can You See Me? Exploring Co-Visiting Between Physical and Virtual Visitors
Areti Galani and Matthew Chalmers, University of Glasgow, Department of Computing Science, UK 31

The Museum Wearable: real-time sensor-driven understanding of visitors' interests for personalized visually-augmented museum experiences
Flavia Sparacino, MIT Media Lab, USA 41

Foundations

Now That We've Found the 'Hidden Web,' What Can We Do With It? The Illinois Open Archives Initiative Metadata Harvesting Experience
Timothy W. Cole, Joanne Kaczmarek, Paul F. Marty, Christopher J. Prom, Beth Sandore, and Sarah Shreeves, University of Illinois at Urbana-Champaign, USA 63

Combining the CIDOC CRM and MPEG-7 to Describe Multimedia in Museums
Jane Hunter, DSTC Pty Ltd, University of Queensland, Australia 73

Today's Authoring Tools for Tomorrow's Semantic Web
Andy Dingley, Codesmiths, Bristol and Paul Shabajee, Graduate School of Education & Institute for Learning and Research Technology, University of Bristol, UK 85

The Virtual Ramp to the Equivalent Experience in the Virtual Museum: Accessibility to Museums on the Web
Liddy Nevile, Motile Research, Australia and Charles McCathieNevile, W3C, France 93

Adding Value to Large Multimedia Collections Through Annotation Technologies and Tools: Serving Communities of Interest
Paul Shabajee, Graduate School of Education and the Institute for Learning and Research Technology (ILRT), University of Bristol, Libby Miller, Institute for Learning and Research Technology (ILRT), University of Bristol, and Andy Dingley, Codesmiths, UK 101

PGP "Pretty Good Practice"

Content Management for a Content-Rich Website
Nik Honeysett, J. Paul Getty Trust, Los Angeles, USA 115

Here and There: Managing Multiply-Purposed Digital Assets on the Duyfken Web site
Marjolein Towler, CONSULTAS Pty Ltd., Valerie Hobbs and Diarmuid Pigott, Murdoch University, Australia 123

Pyramid Power: A Train-the-Trainer Model to Increase Teacher Usage of the ArtsConnectEd On-line Resource
Scott Sayre, The Minneapolis Institute of Arts, and Kris Wetterlund, Independent Educational Writer, ArtsConnectEd Lead Trainer 137

Digital Primary Source Materials in the Classroom
Nuala Bennett, and Brenda Trofanenko, University of Illinois Urbana-Champaign, USA 149

Statistics, Structures & Satisfied Customers: Using Web Log Data to Improve Site Performance
Darren Peacock, National Museum of Australia, Australia 157

How Do You Like To Learn? Comparing User Preferences and Visit Length of Educational Web Sites
David T. Schaller and Steven Allison-Bunnell, Educational Web Adventures, Minda Borun and Margaret B. Chambers, Museum Solutions, USA 167

Evaluating The Features of Museum Web Sites (The Bologna Report)
Nicoletta Di Blas, HOC-DEI, Politecnico di Milano, Maria Pia Guermandi, IBC, Istituto Beni Culturali, Emilia Romagna, Carolina Orsini, Università di Bologna, Paolo Paolini, Politecnico di Milano, Italy 179

Interpenetration

Towards Tangible Virtualities: Tangialities
Slavko Milekic, M.D., PhD, The University of the Arts, USA 189

Making It Realtime: Exploring the Use of Optimized Realtime Environments for Historical Simulation and Education
Chris Calef, Turlif Vilbrandt, Mythworks, USA/Japan, Carl Vilbrandt, University of Aizu, Japan, Janet Goodwin, Aizu History Project, USA/Japan, James Goodwin, University of California, USA 197

Networked Multi-sensory Experiences: Beyond Browsers on the Web and in the Museum
Fabian Wagmister and Jeff Burke, HyperMedia Studio, School of Theater, Film and Television, University of California, Los Angeles, USA 205

Repositioning the Museum

Systematically Speaking: How Do Natural History Museum Web Sites Represent Science?
Roy Hawkey, The Natural History Museum, London, UK 219

Hacking Culture
Pia Vigh, Director CultureNet Denmark, Denmark 231

Storytelling and the Web in South African Museums
Katherine J. Goodnow and Yngvar Natland, International Museums Programme, University of Bergen, Norway, in collaboration with Iziko Museums of Cape Town, South Africa 239

About the Authors

Authors' Biographies 249

Acknowledgements

Our sincere thanks go to all the authors of papers for Museums and the Web 2002, who once again proved that it is possible to write your paper *before* the conference. This rigor helps to document our activities, and contributes to our growing understanding of the intersection of Museums and the Web.

The Museums and the Web site at http://www.archimuse.com/mw2002 is maintained by Adam Brin of Archives & Museum Informatics. Please visit it to find many more papers, live links to sites described, and full descriptions of the demonstrations and workshops presented at MW2002 and all previous Museums and the Web conferences.

This volume was copy-edited by Sheila Trant.

The Museums and the Web conference was managed by Helen Goodman of Event Resources, Pittsburgh, PA.

Introduction

Cyberspace in Our Space

Introduction

Since the Museums and the Web Conference began in 1997, it has examined the confluence between museums, as social institutions with a certain historical tradition and mission, and the "Web", a set of rapidly evolving technologies that combined image, text and multimedia hyperlinking with Internet addressing. Already it is clear that these forces affect each other. Museums have not gone "on" the Web simply to recreate themselves in cyberspace or to become merged with other virtual information stores. Instead, they are beginning to appropriate the technologies of the Web. Museum professionals have engaged the Web as a venue for the exploration of issues they face daily, and increasingly they have found in the Web tools that are of use within the museum. The Web is a becoming part of our physical space just as the museum is re-creating itself in cyberspace; the dialogue between the two is reshaping our sense of museums and museology.

This volume of papers presented at the Museums and the Web 2002 Conference explores the many faces of museums on the Web and the new face of the Web in the museum. The authors probe the many roles of the museum as a social institution, as a series of programs executed in physical space and time, as a collection of objects and knowledge, or as a node in the Web.

Insofar as it proves possible in the coming decades to deploy emerging Web standards to enable museum-like mediation between objects and experience, interpretations of meaning and sources of knowledge, museums will be significant players in the semantic Web. But as the papers in this volume demonstrate, the path to interoperability will first require syntactic agreements and adherence to structural standards for metadata harvesting declarations. These foundations for interoperability are the precondition for broadly available toolsets and newly developing conventions of interaction. The access such standards enable – for the handicapped, for scholars, for the general public and for two-way communication between creators and participants – is being explored. But the underlying assumptions that Web standards will in fact support semantic interoperability and that museums will achieve a mature respect for standards have yet to be proven. Both assumptions are at this point tenuous: museums have been more willing to "adapt" standards than to "adopt" them, and efforts to create structures for universal meaning, or even methods to find generalizable meaning in local perspectives, have so far proven elusive.

Whether knowledge of our heritage is ultimately integrated into a "world wide" Web of knowledge, or not, museums have embraced the Web and are mastering the management methods, pedagogical techniques, and assessment strategies needed to support effective institutional programming. Leading museums have introduced structures, technologies and business relations from the world of commercial organizations into museum "content management", and have professionalized their Web operations. Where home-grown programs once reigned, professionally planned, inter-institutional, technologically sophisticated Web-based programs raise a challenge. Most importantly, a rigor in evaluation rarely seen in the assessment of non-technologically mediated museum programming is being

BEST COPY AVAILABLE

introduced, and indeed, expected, in Web programs. Improved methodologies for evaluation in the initial planning, pre-testing and post-implementation phases are being implemented regularly. Museum management is coming to expect an explanation of the relative effectiveness of different means of interpretation. In a relatively short time, Web-based feedback mechanisms have helped museums fine-tune the instruments of their communication using sophisticated evaluation tools.

Museums' successes derive increasingly from integration of the Web into the museum. Not only is this affecting work processes such as content management, education and evaluation; but on-site exhibitions and interpretation are also beginning to benefit from the Web-based delivery of context for objects on exhibition. The Web as a potential threat to the on-site, physical, real world museum is confronted, by moving the virtual exhibition – located somewhere else, in cyberspace, and most often approached from somewhere other than the museum homepage – closer to the physical space; in-gallery interactives have begun to bridge the gap between the virtual exhibition – a separate venue requiring distinctive content, and competing for funds and staff energy– and ongoing museum exhibition programming, with all the strengths of the museum's identity and physical presence. Web-driven handhelds married to wireless communications facilitate interpenetration of the virtual and the real. This marriage of physical and conceptual space is inspiring experiments, especially by university-based researchers, and generating a new social and intellectual construct: mixed reality. While it is too early to know how people will ultimately integrate mixed reality into their everyday life, museums and academic researchers working with museums are modeling the options.

Museums have adopted the Web in as many different ways as any other social institutions in our society. Already there are some signs that the museum as an institution may be significantly changed by the encounter. Just as commerce is not fundamentally affected by B2B applications (which make commercial transactions more effective) but might be fundamentally altered by reverting to on-demand custom fabrication (as with Dell Computers) or the social exchange of free information (as with Amazon), so also museums could be fundamentally altered if other types of institutions usurp their role as custodians and gatekeepers of art, impinge on their position as collectors and interpreters culture, or offer methods that make possible the aggregation of information from many sources. Or the Web could provide tools for museums to reconnect with their communities, reasserting their relevance by empowering their visitors, virtual and physical.

The Edge of the Web

This year, the edge of the Web is a wedge that penetrates museum galleries; wireless computing and handheld devices have moved the Web from cyberspace into museum space. The first section of the *Museums and the Web 2002 Proceedings* reports on several projects exploring the potential for innovative interpretation of museum objects and ideas, delivering networked digital objects to people using handheld devices in wireless networked spaces and watching as those people communicate with each other, with the "smart space" around them and with linked knowledge sources.

Robert Semper from the Exploratorium and Mirjana Spasojevic of Hewlett-Packard Laboratories introduce the state of the art in the "electronic guidebook", a metaphor and a physical, portable device that uses a wireless Web-based network to "extend the museum experience". As these devices are becoming smaller, more affordable and more capable, museums will soon be able to count on them to facilitate delivery of complex information to visitors. What is not yet clear – and is the subject of this and other research – is what kind of information should be delivered, to whom and at what time. We still have much to learn about how this can be done in a way that attracts attention to the artefact or idea we are communicating, and enables communication among visitors rather than inhibits it. There have been several attempts to implement full-blown hand-held guides within galleries (including those at SFMOMA reported by Peter Samis at the ichim 2001 meeting in Milan, Italy); Semper and Spasojevic step back to examine requirements for success in a systematic fashion, considering infrastructure, interface and content development. Two problems are identified in this research: 1) the potential for interference with the social experience of visiting, as the device gets in the way of communication between visitors, and 2) actual physical interference, as holding the device impairs visitors' ability to interact with the exhibit.

At PARC, a group of researchers specifically addressed the social issues identified by Semper and Spasojevic, exploring the relative advantages and disadvantages for social interaction of different methods of delivering audio in a gallery tour environment, including one where visitors shared similar aural experiences. Their conclusions point both to new strategies for technological mediation and to the need for further studies of computer-human interface options and effects. At the University of Glasgow, Galani and Chalmers conducted a systematic analysis of the social context of the museum visit, exploring how it can be affected by technology. They studied ways in which different visitors to an exhibit actually interact together once they are there, with each other, the gallery, the Web and the virtual environment. Their results similarly point to technologies that could substantially overcome the barriers to social interaction in a Web-facilitated visit, and could possibly integrate Web and physical visits.

The second barrier is physical. In her exploration of the past and potential of the "Museum Wearable", Flavia Sparacino of the MIT Media Lab examines ways in which wearable computers could augment the gallery environment, enhance visitors' experiences, and in the future even enable visitors to use their hands freely.

Foundations

While museums may have been conceived as independent sources of interpreted knowledge, the Web has pointed out the weaknesses of this isolation for the construction of complex understanding. It takes little advantage of the inter-relationships between and among disparately located museum objects. Together museums' collective knowledge can only be identified, navigated, explored, and integrated if its structure is explicitly declared. Interoperable knowledge models could drive tools to enable access across barriers of differing abilities, disciplinary perspectives, and documentation traditions. But the levels of standardization that museums will need to adopt, minimally to move data successfully from place to place (without carrying much meaning), to say nothing of the inherent complexity of the project of moving information created in one intellectual perspective into meaning in another context, are daunting. There is much scope for research in this area, and the papers in the "Foundations" section touch on significant areas of standards definition and application, in metadata harvesting, data modeling, the semantic web, accessibility, and annotation.

The metadata harvesting experiments conducted by the multi-departmental research team at the University of Illinois begin to peel back the layers of standards infrastructure required to easily integrate information from disparate technical environments, media formats, and disciplinary perspectives, and to deliver it to users with different capabilities, knowledge and tools. In the "Illinois Open Archives Initiative Metadata Harvesting Experience", supporting the OAI server and harvesting technology per se is not even seen as a fundamental barrier most museums couldn't overcome. Attention is focused on the next level of the problem – that of standardizing the content of metadata, which is shown to be quite difficult even when a very simple common denominator, the Dublin Core, is all that is attempted.

Jane Hunter tries to bridge a much more complex set of metadata models for knowledge representation in mapping the CIDOC Conceptual Reference Model (CRM) designed to describe museum data to MPEG-7 designed to describe multimedia, in order to describe multimedia museum content. Hunter uses the Resource Description Framework (RDF) in order to express the enhancements needed to the CIDOC Model to incorporate critical MPEG-7 concepts. Andy Dingley and Paul Shabajee pursue the theme of content integration as they envision how today's authoring tools, including XML and RDF, can be used to construct the "semantic Web" envisioned by the W3C. Using the ARKive project as their test-bed, they are prototyping the interoperable knowledge environment of the future. Liddy Nevile and Charles McCathieNevile take the dialog a step further as they introduce a series of W3C standards designed not simply to enable users of the Web to move across data sources, but also to enable access in different modalities and with different enhanced interfaces for users with disabilities.

An alternative approach to building a semantic network is explored by the team at the University of Bristol in their examination of the use of annotation to enable the collaborative indexing of a common data store according to the perspectives of various communities of interest. Using a relatively low technology approach in place of automated tools and highly sophisticated mechanisms for translating between schemas, they propose a cooperative, social means to bridge intellectual perspectives and enable cross-disciplinary access.

PGP (Pretty Good Practice)

After less than a decade, it is evident that museum new media practice for the Web has introduced new rigor throughout the process of managing content, delivering learning experiences and evaluating results. Indeed, we now have "pretty good practice" – if not "best practice" (something we will probably never achieve and might not recognize if we did). Like pretty good privacy, it's the best we've got.

Two papers in this section address the need that museums, and other organizations with lots of data, have to manage their "assets" using sophisticated tools that make possible the maintenance of lively, timely, and rapidly changing Web sites. The methods introduced by the Getty, to integrate data from its museum and many research centers, and by the Duyfken project, to bring content from numerous sources together on an hourly basis, are indicative of the types of content management practices that are becoming necessary everywhere.

The purpose, of course, is to create sites that have the depth to be attractive and the range of digital objects to appeal to a variety of different visitors – what is referred to elsewhere in this volume as having "a low threshold and a high ceiling" (Nevile and McCathieNevile). Pretty good practice at this point is to develop programs that bring educators in the classroom and in the museum together to develop strategies for enabling and improving use. In "Pyramid Power", Scott Sayre and Kris Wetterlund report methods of teacher training used to introduce teachers to museum resources and explore how they can make good use of them. The power of their model comes in the 'train the trainer' cascade. In their exploration of issues surrounding the integration of digital primary resources in the classroom, Bennett and Trofanenko describe a collaboration among teachers and a number of cultural heritage organizations.

In the end, the question is whether the investments we are making, in teaching tools and in rich museum Web sites, really pay off. Darren Peacock proposes a framework and suggests quantitative measures from Web log data to assess user satisfaction with the Web site of the National Museum of Australia. A group of commercial designers and museum educators turn to "Comparing User Preferences and Visit Length of Educational Web Sites" in a qualitative study that bridges learning theory and interactive media development. And a group of Italian academic software engineers apply their design-based metrics to the assessment of some well-known museum sites with results that are informative for anyone planning an evaluation or designing a site.

Interpenetration

As the edge of the Web insinuates itself into museum space, university researchers are developing new tools and techniques to wrap cultural content in digital context and present it in an engaging and intuitive manner. The results challenge the boundary between the virtual and the real, and ask us to consider how we might best engage with culture, digitally.

Physician and experimental psychologist Slavko Milekic is defining a new domain – Tangiality – that bridges Reality and Virtuality. Here museum data and museum visitors may soon come together through haptic interfaces. At the University of Aizu, researchers are exploring the use of tools and metaphors developed for games to make visualizations of historical simulated environments for museum interpretation. And at UCLA, in the Hypermedia Studio of the School of Theater, Film and Television, installation artists are constructing environments where performers, in real and virtual spaces, interact with audiences that are vital participants. In each of these mixed reality environments, real people, digital input and data interact in an environment, respond, and adapt. They suggest that the interpenetrated world of mixed-reality is, increasingly, the world that museum programming and the use of the Web will exploit.

Repositioning the Museum

So, will this change the museum? Has it changed the museum? Communications and media theory predict that the introduction of new media could fundamentally alter the museum message. Indeed, after a relatively few years in the life of a fairly conservative type of institution, it is becoming apparent that the Web can affect the fundamental nature of museums. The Web is not just a brochure to advertise the museum (as it was in the first couple of years), or an alternative venue in cyberspace in which to implement museum programs (as it has been for the past couple of years). The final section of this book offers tantalizing hints about the impact the Web may have on the nature of the museum as a social institution, its mission and purpose. If, as Pia Vigh suggests vis à vis net art in "Hacking Culture", the medium is the message and the message can only be conveyed in a new social context beyond the museum, then the Web will have proved a serious challenge to museums. But if, as Roy Hawkey suggests, the medium makes possible the communication of a message that had been subjugated to a minor position in traditional museums where it failed to fit into the medium of exhibition display, then the Web could liberate the museum to realize science education in a way that was previously constrained. Or perhaps, as in the culturally sensitive exploration of potentialities of Web interactions with museum missions elaborated by Goodnow and Natland, the museum can realize a new mission by grappling with precisely this complex interaction between what a museum represents as an institution and what museology might enable.

And remember ...

This volume is entitled *Selected Papers* from Museums and the Web 2002. We are delighted that it reflects the rich discourse that characterizes the conference and this community, but as the choice of which conference papers to print becomes ever more difficult, we remind readers that four times as many papers from the conference are included on the CD-ROM that accompanies this book. And even the CD does not fully represent the meeting. To appreciate the entire range of creative discussion that accompanies this annual encounter between the Museum and the Web, it would be necessary not just to review all the papers, but also to experience the 14 pre-conference workshops, 50 demonstrations, 35 exhibits and 15 mini-workshops, as well as the performances, Crit Rooms and Usability lab - interactive events – which are not documented on-line or on the CD. Ultimately, a conference is a social event: the experience cannot be fully represented in any publication. Nevertheless, we hope this volume stimulates more research and experimentation, and yet more publication about Museums and the Web.

David Bearman and Jennifer Trant
Archives & Museum Informatics
http://www.archimuse.com

The Edge of the Web

The Electronic Guidebook: Using Portable Devices and a Wireless Web-based Network to Extend the Museum Experience

Robert Semper, Exploratorium, and Mirjana Spasojevic, Hewlett-Packard Labs, USA

Abstract

Recent advances in wireless network technologies create the potential to significantly enhance the experience of a visit to a museum. On the exhibit floor, visitors carrying wirelessly connected portable devices can be given opportunities for exploration, sharing, explanations, context, background, analytical tools, and suggestions for related experiences. When these devices are part of a Web-based network, they can help extend the museum visit: in advance, through activities that orient visitors; and afterward, through opportunities to reflect and explore related ideas.

The Electronic Guidebook project is a study of visitors equipped with such technologies, conducted by the Exploratorium in partnership with researchers at Hewlett-Packard Labs and the Concord Consortium. The project is investigating how a Web-based computing infrastructure can provide museum visitors with an augmented museum experience so that they can better plan their visit, get the most out of it while they are in the museum, and be able to refer back to their visit once they have returned to their home or classroom. The goals are to understand what technological infrastructure supports this extended museum experience, and to obtain preliminary data on how different aspects of the technologies, and the content delivered through them, affects engagement with the exhibits and pre- or post-visit learning activities.

Keywords: handheld computers, wireless networks, electronic guides, mobile computing devices

Background

The project created a test-bed and tested a network using a variety of handheld computers and radio-frequency identification tags to link visitors with exhibit-related content delivered by a Web-based server. Visitors in the study were able to access Web-based content, including text, images, video, and audio, during a visit. In addition, they were able to construct a record of their visit by bookmarking exhibit content, creating images, notations, and other artifacts, and to access that record on a personal Web page in the museum or following their visit.

Rationale

The exhibit environment in museums excels in providing the public with direct experiences with phenomena, ideas, and objects. But while the exhibits provide compelling opportunities to stimulate inquiry and exploration, they are not, by themselves, completely successful in supporting conceptual learning, inquiry-skill-building, analytic experiences, or follow-up activities at home or school. An ideal learning experience with exhibits would include ways of capturing the visiting experience for later reflection, being able to access additional material that provides a context for the exhibit, and extending the interaction with the exhibit beyond simple manipulation. If museums are to fulfill their potential role as multi-dimensional educational institutions, they need additional mediation techniques that support these needs. And these techniques need to augment the simple exhibit experience without destroying the informal ambiance that is the hallmark of museums. In effect, what is needed are tools that can provide the kind of contextual support with exhibits that is offered by the familiar travel guides but defined in a new interactive way.

Guides are an important part of our travel experience. We use them in advance to plan our trips. They provide us with an historical context for the country or city we are going to see. We take them along to provide quick access to essential information like addresses and transit information. We use them as in-depth guides to tours of historical land-

BEST COPY AVAILABLE

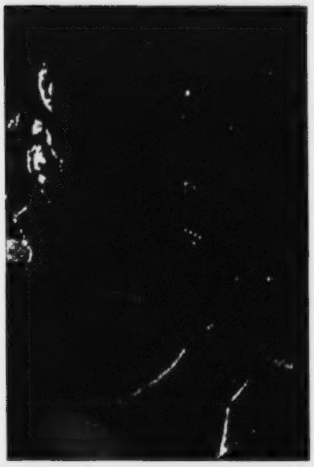

Fig. 1: Visitors using the Electronic Guidebook in the Exploratorium

marks. They help us translate words and exchange currency during our visit. We personalize them with notes, hotel brochures and postcards from the places we have stayed. We refer to them after the trip to remind us of our visit, and they form the basis of talking about our trip with our friends. The best guidebooks become dog-eared from use by the end of a trip and become keepsake reference points for whenever we "revisit" the journey. In short, they serve as an ongoing tool for enhancing our direct visit experience with a place and culture.

Emerging computer and network technologies have the potential to significantly enhance the experience of a visit to a museum, giving visitors the opportunity to transcend the traditional limits of time and physical space as they engage in a museum's offerings. Electronic guides, enhanced via interactive technology, can not only serve as reference information, but also provide expert guidance, dynamic advice, recommendations for further inquiry, and other learner supports no paper-based travel guide can provide. On the exhibit floor, additional information, guidance, or feedback at the right time can turn a simple interaction with an exhibit into a meaningful learning experience. Inexpensive portable devices can give visitors opportunities for exploration, sharing, explanations, context, background, analytical tools, and suggestions for related experiences.

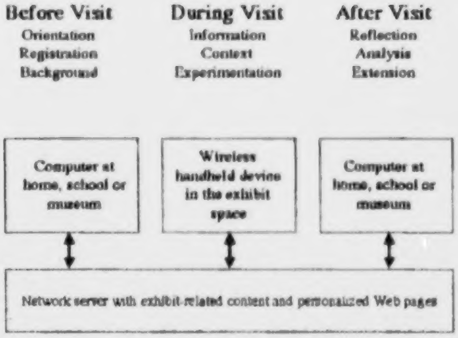
Before Visit
Orientation
Registration
Background
During Visit
Information
Context
Experimentation
After Visit
Reflection
Analysis
Extension
Computer at home, school or museum
Wireless handheld device in the exhibit space
Computer at home, school or museum
Network server with exhibit-related content and personalized Web pages

Fig. 2: The Museum Experience Extended through a Wireless Web-based World

The visit can be extended through sustaining learning opportunities: in advance through activities that orient visitors, and afterward through opportunities to continue reflection and explore related ideas. By adding personalized technology to the museum experience in these ways, we can greatly enhance the visit's educational value.

Research Project

The Electronic Guidebook project (http://www.exploratorium.edu/guidebook/) is a National Science Foundation funded activity designed as a proof of concept research study to explore the potential of portable computing devices and a wireless Web-based network. Headquartered at the Exploratorium (http://www.exploratorium.edu), the project is a partnership of the Exploratorium, Concord Consortium, an educational research and development organization, and HP Labs, the research and development laboratory of Hewlett-Packard Company. The Exploratorium worked with the Concord Consortium to identify the technical issues and potential uses of handheld devices in the museum setting, and the museum formed a partnership with HP Labs to develop and test a Web-based, wireless network deployable in the Exploratorium' physical space. The aim of the project is to investigate how a mobile computing infrastructure can enable museum visitors to create their own "guide" to the Exploratorium, using a personalized, interactive system to better plan their visit, get the most out of it while they are in the museum, and be able to refer back to their visit once they have returned to their home or classroom.

The Electronic Guidebook project focused on three simultaneous strands of investigation:

BEST COPY AVAILABLE

1. information technology infrastructure (networked components delivering the information);
2. human computer interface issues (form factors, interfaces, and usability issues for various audiences); and
3. content development (design, formatting, and topical content of deliverable information).

The project was designed as a proof of concept study to explore potential avenues for future research and development. It was not envisioned to support the implementation of a fully functional system, but rather to point the way for future developments. A key idea was to create the test-bed within the existing museum context and use off-the-shelf technology to the extent possible so as to be able to explore this technology in a realistic setting. Often technology projects are tested in idealized conditions without the distractions and the messiness that occurs in the real world. Our goal was to study the interaction of all of the elements, technology, functionality and content, within an authentic situation.

The project worked with an existing set of museum exhibits that are part of the "Matter World" exhibit area. These exhibits that addressed topics in sound, mechanics and heat are highly interactive, distributed in space, and individual in design. Because these exhibits have been in the museum for many years, there exists a rich set of supplemental materials (text, images, teaching activities, exhibit developer interviews, etc.) already developed. While this material could provide a rich context for the visitor, it is usually not readily accessible because it is scattered in files, offices, and the brains of staff members.

The Test-bed

Initial developmental studies which tested different devices, networks and servers were pursued in 2000 by Concord Consortium and the Exploratorium. For the testbed we augmented the museum environment with technologies developed by the Cooltown research program at HP Labs (http://cooltown.hp.com) [Kindberg (2001); Kindberg (2000)]. In the Cooltown scenario, all physical entities (people, places and things) have a 'Web presence' [Debaty (2000); Kindberg (2002)]. Nomadic users navigate from the physical to the virtual world by picking up links to Web resources using a variety of sensing technologies such as infrared receivers and barcode readers. Those readers are integrated with their handheld device, which is wirelessly networked. Users access electronic services by using handheld devices to pick up URLs from barcodes or infrared 'beacons' attached on or near the objects of interest. Those services are provided using wireless Web technology: Web browsers on the handheld devices and Web servers in the environment. The services can be adapted for users based on their context, e.g. their identity, location, device capabilities, personal interests and preferences.

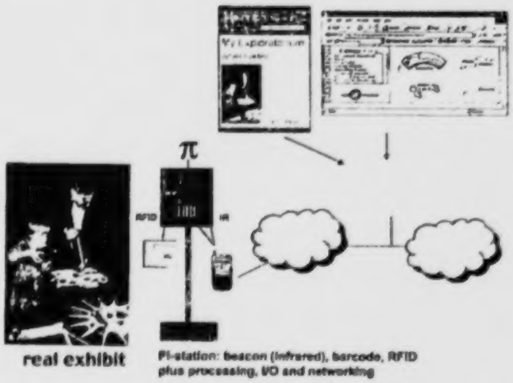
π
RFID
IR
real exhibit
Pi-station: beacon (infrared), barcode, RFID plus processing, I/O and networking

Fig 3: Setting: The Augmented Interactive Museum)

The Electronic Guidebook project test-bed consists of a wireless handheld carried by the visitor, a "pi station" (point-of-information station) at each exhibit, and a content server connected to the World Wide Web on the Internet and accessible both inside the museum and from home or school. The visitor located in the museum space was identified using the handheld by a short range IR transmission, an RFID sensor (http://www.aimglobal.org/techologies/rfid/), or a barcode swipe. Data transmission to the handheld was provided using a radio LAN system in the exhibit space that provided connectivity to a Web server which is connected to the public Internet.

Each pi station is a tall stand with a base that can contain a laptop and a panel at eye level that can contain an IR beacon (http://cooltown.hp.com/beacon_full.htm), RFID reader, barcode, and/or camera. At each pi station, the visitor can "sense" the URL of the main page for the corresponding exhibit. The

BEST COPY AVAILABLE

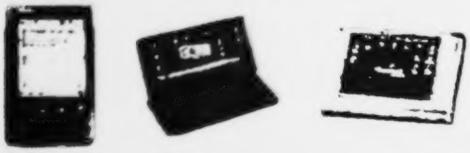

Fig. 4: Handheld devices L to R HP Jornada 548, HP Jornada 690, Hitachi ePlate)

page is then retrieved from the Web server and displayed with a browser on the handheld. The pi-station was designed to help test different interaction modalities and to facilitate working with different exhibits for different tests. In future system design, some of this functionality might be built into the exhibit itself.

Three types of hand-helds have been used for the studies: HP Jornada 690's and 720's (a Hand-held PC with a physical keyboard), HP Jornada 540's (a Pocket PC with on-screen keyboard), and Hitachi ePlates (touch-screen tablets). The hand-helds are connected to the network using 802.11b cards which provide wireless connectivity at 11 mbps. A single 802.11b base station is connected to the Exploratorium's network near the exhibits.

The Jornadas have IR communication ports and Esquirt software that allows them to collect URLs from Cooltown IR beacons. When the visitor points the Jornada at an IR beacon, the beacon transmits to the Jornada the URL of the on-line content associated with that beacon. A barcode reader is attached to the Hitachi ePlate. The raw ID provided by the barcode is mapped to a URL, which is then treated like the URL's provided by beacons.

Information and contextual material for each exhibit are displayed on the hand-helds using a modified Web browser. When the visitor picks up a beacon or barcode for an exhibit, the main page for that exhibit is automatically displayed. The browser's user interface has been simplified so that the only browser buttons are forward, back, home, and a bookmark, which we added to allow visitors to record visited pages in their personal scrapbook.

Using the System

Visitors carry the electronic guidebooks with them as they walk through the museum floor. As they walk up to a particular exhibit, the guidebook provides interaction, information, and ideas about the exhibit. When approaching the Exploratorium's Echo Tube exhibit, for example, visitors see a small Web page on their handheld guidebook, suggesting ways of interacting with the tube, such as clapping and counting to measure the speed of sound. The guidebook also provides visitors with intriguing questions to consider (such as, "What does a dog bark in the tube sound like?") and a way for them to enter their own questions and observations (which other visitors can then access). In addition, visitors can learn more about the science of echoes (with deeper explanations than available on the exhibit label), real-world connections (for example, recommended locations in the San Francisco area for hearing echoes), and the history and evolution of the exhibit itself (including stories from the exhibit developers). The system allows visitors to save items of interest to a personalized Web page that they can access after their visit at home or school.

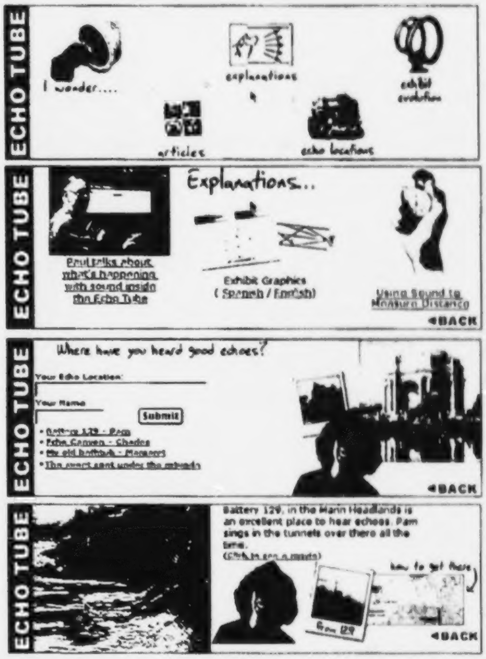

Battery 129, in the Marin Headlands is an excellent place to hear echoes. Pam sings in the tunnels over there all the time.

Fig. 5: Echo Tube Pages

All exhibits have a main page whose URL is stored in the pi station, as well as 3-4 sub-pages. When the main page gets downloaded by the browser (after picking up a URL through a beacon or a barcode),

BEST COPY AVAILABLE

visitors can click on any part of the image to get to the "nuggets" page which provide links to several areas for further exploration of content related to the exhibits.

In addition to the individual exhibit pages, visitors can access the Electronic Guidebook home page. The home page contains a list of exhibits and a map of exhibit locations. It also gives visitors access to a personal scrapbook page on the server where they can record URLs of interesting material which can be reviewed later on the WWW from any browser.

Research Studies

The research team that is studying the use of the test-bed is made up of museum educators, scientists, educational technology researchers, learning researchers and human-computer interface researchers. The project started with a set of questions that the test-bed was designed to help answer.

Technology Development

- What is the feasibility and effectiveness of various wireless receiving and transmitting solutions for an exhibit space including distributed and/or centralized infrared and radio transceivers?
- What is the appropriate distributed computing power relationship between the handheld device, the exhibit computer, the transceiver system and the network servers?
- What is the optimal integration of the dynamic user database in the Web environment?
- What are the appropriate handheld device data sensors for the museum exhibits and the design issues surrounding the interface between these sensors and the exhibits themselves?

Interface Design

- What is a feasible system for providing handheld devices to the public? How can we distribute these devices equitably and to a diverse cross-section of our visitors? Will the devices be used by individuals or by multiple visitors in a group? How will security be handled? What are the logistics of carrying them on the exhibit floor?
- What are the issues surrounding the use of these physical devices during a visit to an exhibit environment? How can we integrate the devices into the complex social activity of a museum visit without fragmenting a group? How can they be comfortably handled during a prolonged visit?
- What is an appropriate software, hardware and exhibit graphics interface design for the experience at the exhibit, at home before or after a visit and in the Learning Studio as part of the museum visit? What are the important accessibility issues? For example, how many visitors have Web access? For those who do not have Web access at home, in what ways can we provide them with closure to their experience?
- What are the key hardware and software design issues (screen size, speed of use, amount of material) that form a threshold of feasibility for this idea? What are the specific design elements (text, graphics, etc) which are important for the experience? What are the specific mediation components that this technology can support at exhibits (capturing information, providing context, creating representations, taking data)?

Educational Potential

- What is the effectiveness of the Electronic GuideBook system for capturing experience, providing context and extended interaction? How can we create resources that stimulate exploration and reflection rather than having the handheld just seem like an electronic textbook?
- What kinds of learning can we facilitate: Cognitive? Affective? Attitudinal? Skill-based? What are the key issues that might motivate the use of these devices? How can we create resources that stimulate exploration and reflection?
- What is the value of the different project components to the overall educational experience (i.e. pre visit, at the exhibit, in the museum's classroom spaces, in the home after the visit)?

The initial work has been focused on developing the test-bed and initiating the study with a series of research questions. These studies have focused on the following ideas:

- What are the basic affordances? For example, is the device easy to carry? Are screen graphics easy to interpret? Is the wireless connectivity reliable? Is the 1-2 meter range for reading a Cooltown beacon more convenient for users than the 1-10cm range for scanning a barcode?

- What is the visitors' attention to artifacts? That is, how much attention do the users pay to each exhibit or piece of on-line content, and what is the quality of this attention? For example, did the users look at a particular page? If so, did they glance at it or read some text on it, for example? Were the users mainly paying attention to the exhibit, paying attention to the content on their handheld device, or fiddling with the device itself (e.g. trying to solve some problem with the interface)?

- What are the visitors' paths through physical and virtual space? What is the sequence of points in physical space (exhibits) and virtual space (Web pages) through which the users pass as they visit the museum? Of particular interest are trajectories that correspond to high-level user events such as following a virtual link to get more information on a topic, deciding to walk over to another exhibit, shifting attention back and forth between the physical and virtual at a single exhibit.

The next set of studies (currently underway) are designed to help us better understand how different types of content (including different multimedia formats, types of suggested activities, and the ability for visitors to add their own observations and ideas) affect users' experience at the Exploratorium and in what ways mobile electronic resources support or hinder their experience in the museum context. Also planned are studies of the "out of museum" use of the network to explore its utility before and after a visit.

Overview of Studies

To date, we have carried out several rounds of informal visitor studies. The goal of these tests was to get a general sense of how visitors reacted to the system and to uncover major issues. The users for these initial tests were 12 local school teachers (9 female, 3 male), 14 other adults (8 male, 6 female), and 8 young people, ages 10-17, (5 male, 2 female). The adults ranged in age from about 25 to about 50. All users were fluent speakers of English with no major physical disabilities.

The evaluation approach selected for the tests was participant observation with a project member shadowing a visitor or a small group of visitors. Visitors were directed up to the general area of the instrumented exhibits. They interacted with those exhibits and sometimes nearby non-instrumented exhibits. The project members observed the visitors' actions and reactions, and helped them out when they had major problems. Semi-structured interviews were held afterwards. The tests were done when the museum was open to other visitors.

For these initial tests, we used HP Jornada 690's, except for one subject who used an Hitachi ePlate. In-depth content has been developed for six exhibits, including one pair of exhibits that are thematically related to one another. Most of the pages have heavily graphical content, with a few having mixed text/graphics content and parallel audio content. The content is formatted for the Handheld PC display size. A stand was placed next to each exhibit, with a barcode and IR beacon containing the appropriate URL for each exhibit's content. Content was designed primarily for an adult audience.

Lessons Learned

The overall reaction from the visitors was positive. They enjoyed reading the on-line content. Some adults felt that the content would not be interesting to children. However, children who used the system did actually spend a fair amount of time reading the on-line content and trying out possibilities suggested by it. Visitors frequently commented that the hand-helds were really fun, novel devices. On the one hand, this helped keep them interested. On the other hand, the "wow" factor of using the device distracted them from paying attention to both the exhibits and the on-line content.

Above all, visitors liked the idea of being able to bookmark information to look at later. Both teachers and children thought this feature would allow the children to play more during their museum visit, completing related homework assignments after the visit.

Aside from technical problems discussed below, the main negative comments were that using the handheld tended to distract people from playing with the exhibits. The issues here are both mechanical and cognitive. Holding the handheld prevents people from freely using their hands to manipulate the exhibits. Moreover, reading the content is itself distracting from interacting with the exhibit.

Using hand-helds in a museum setting may also disrupt normal social interactions between members of social groups, because each visitor can get lost in the world of the device and pay less attention to the rest of his group. For example, we observed two boys moving around the museum floor together, talking about the exhibits and working things out together. Yet when we gave them each a handheld, one spent more time reading the information on the device while the other was still interested in seeing as much as he could. On some occasions, when they needed two people to make an exhibit work, they had to spend time trying to find each other again. We believe that these issues can be overcome by designing a system to specifically support visitors' conversation patterns, such as the system tested at an historic house [Aoki (2000); Woodruff (2001)] where the electronic guidebook became a third party in the visitor conversation.

Handheld design and features

Hand-helds do not routinely come with straps. Carrying cases are only available for certain models. This represented a serious problem in the Exploratorium because visitors must have one or both hands free to operate the exhibits. When early trials made it apparent that a strap was necessary, we had to resort to tying a strap around the Jornada 690's hinge. However, even with a strap attached to it, a Hitachi ePlate was too heavy to use for a prolonged time during the visit.

Our original plan involved comparing audio renditions of content to parallel graphic versions, as in a recent handheld project at an historic house (Aoki 2000 and Woodruff 2001) This study found that visitors preferred audio content to visual content, presumably because it left their eyes free to look at the exhibits. However, the Exploratorium is extremely noisy. When the Jornada is held in the hand, sound from its speakers is essentially inaudible. Some visitors were not even aware that any sound was playing. Headphones can be used, but the Jornada 690 happens to lack a headphone jack (a problem fixed in the Jornada 720). Therefore, the visitors had to be prompted, by on-screen graphics, to hold the hand-held up to their ears.

Robustness issues are a concern for use in an interactive science museum. Robust devices are essential if on-line content is to be seamlessly integrated into everyday activities. The hand-helds struck most visitors as too fragile to withstand being dropped or bumped. Many exhibits at the Exploratorium involve materials like sand and water, which handhelds were not designed to withstand.

The wireless network performance was adequate for downloading Web pages and short audio files. However, network performance can be disrupted by obstacles such as elevator shafts. It is essential to test network performance in the place where it will be used and to check carefully for possible "holes" where performance is bad. Multiple wireless base-stations may be required in such a case.

Picking up Beacons and Barcodes

We found that visitors were able to sense beacons or barcodes with the handheld devices with little trouble. Other Cooltown demos have used a beep to indicate successful beacon or barcode pickup. Since this sound cannot be heard in the Exploratorium, visual feedback must be provided instead.

Likewise, we found that a brief hands-on introduction is required: sending visitors off with only verbal instruction isn't sufficient. However, all visitors became fluent with the devices quickly. The one person who used a barcode reader behaved similarly to those using beacons.

The visitors never accidentally picked up an IR beacon for a different exhibit. However, some visitors accidentally picked up an exhibit's beacon repeatedly, because they happened to hold the Jornada so that its IR port pointed at the beacon while they were browsing the on-line content. We have since added features to the browser to minimize duplicate pickups by ignoring them if they occur very soon after original pickup, and by querying the visitor about whether to redisplay the exhibit top-level page if the delay is longer.

Graphical User Interface

In these studies, we gave visitors only extremely brief instructions on how to use the browser interface. Visitors familiar with hand-helds and browsers had no trouble using the interface. However, some visitors inexperienced with hand-helds had significant problems successfully selecting objects on the touch screen, e.g. they dragged the stylus along the surface rather than tapping. The visitors also weren't familiar with standard browser buttons such as " back". However, these visitors did use the system successfully after brief hands-on instruction. For the last two trials, we added a practice exhibit to the check-out process to alleviate this problem.

Content Design and Navigation

Visitors were navigating the content using a stylus and clicking on images or text hyperlinks presented on the Web pages. When the top-level navigation page for an exhibit consisted of one large picture (e.g. Fig. 6), some visitors had trouble figuring out that they could click on part of it or which parts of it were click-able. Some visitors expressed a desire to be able to see how much content there was for each exhibit and which parts of it they had/hadn't seen so far.

Many visitors to the Exploratorium seem to migrate from one exhibit to the next interesting-looking one. In such cases, the visitors only need to be able to locate a pi station once they have already approached the chosen exhibit. This is not difficult so long as the pi station has a reasonably distinctive appearance. However, other visitors may wish to follow a tour or see a group of thematically related exhibits. Or the on-line content may suggest another exhibit closely related to the one they are at. It is unclear how to best direct visitors from one exhibit to another exhibit. Some of our test visitors had difficulty locating pi stations on the same floor and only a modest distance away, even when marked by bright orange flags. The on-line map wasn't used much - probably because it was not prominent enough in the on-line interface or was too difficult to relate to the 3D environment.

Future Research Plans

We plan to continue our research with the testbed in the future with a specific aim of conducting studies with different audiences (general public, school field trips, Explainers (floor interpretive staff), teachers, members and staff.) An important feature of our next study will be to examine the potential use of the system in before-and-after visit situations. Our aim is to continue to explore issues of using wireless networks and portable computers to expand and extend the museum visit experience.

Conclusion

An additional part of the Electronic Guidebook project is a conscious effort to stimulate discussions within the museum community and between museums and industry about the potential of using these systems in a museum setting. To initiate these discussions, the Exploratorium sponsored the first Electronic Guidebook Forum in the fall of 2001. This two-day forum brought together 39 researchers and

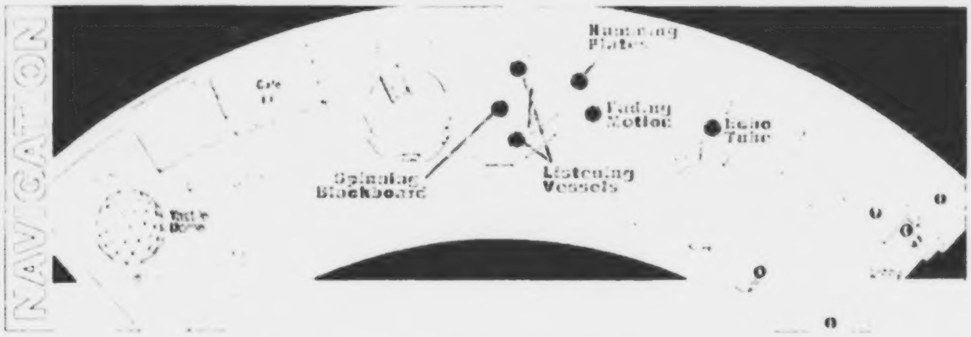

Fig. 6: Electronic Guidebook Home Page

BEST COPY AVAILABLE

developers from industry, academia, and the museum world for discussion of the latest findings on the application of handheld computers and wireless networks in museum exhibitions. The forum discussions centered on interrelated aspects of electronic guidebook projects in museums and on emerging questions from the field. The format included full group discussion of these topics, as well as discussion in small groups on lessons learned and recommended next steps. The goal of the forum was to identify key issues in the museum field that will inform further work on wireless handheld devices and stimulate research and implementation. An in-depth report on the forum is available at http://www.exploratorium.edu/guidebook/.

Acknowledgements

The authors would like to sincerely acknowledge the hard work and support provided by the many people who have contributed to this project. They include:

Katherina Audley, Stephen Bannasch, Doug Conaway, Philippe Debaty, Margaret Fleck, Marcos Frid, Caroline Gattein, Joshua Gutwill-Wise, Ron Hipschmann, Sherry Hsi, Tim Kindberg, Eamonn O'Brien-Strain, Mike Petrich, Rakhi Rajani, Natalie Rusk, Larry Shaw, Robert Tinker, Quan Tran, Adrian Van Allen, Noel Wanner, Karen Wilkinson.

This material is based upon work supported by the National Science Foundation under Grant No. ISE-9901985.

References

Aoki, P. M. (2000) and Woodruff, A. Improving Electronic Guidebook Interfaces Using a Task-Oriented Design Approach. *Proc. 3rd ACM Conf. on Designing Interactive Systems*, New York, NY, Aug. 2000, 319-325.

Debaty, P. (2000) and Caswell, D. Uniform Web Presence Architecture for People, Places, and Things. HPL Technical Report HPL-2000-67, to appear in the IEEE Personal Communication magazine on Pervasive computing.

Kindberg, T. (2002). Implementing physical hyperlinks using ubiquitous identifier resolution. HP Laboratories technical report HPL-2001-95r.1. To appear in Proceedings 11th International World Wide Web Conference, Honolulu, May 2002.

Kindberg, T. (2001) and Barton. J. A Web-based Nomadic Computing System. *Computer Networks*, Elsevier, vol 35, no. 4, March 2001, 443-456.

Kindberg, T. (2000) et al. People, Places, Things: Web Presence for the Real World. Proc. 3rd Annual Wireless and Mobile Computer Systems and Applications, Monterey CA, USA, Dec. 2000, 19.

Woodruff, A. (2001), Aoki, P.M., Hurst, A., and Szymanski, M.H. Electronic Guidebooks and Visitor Attention. Proc. 6th Int'l Cultural Heritage Informatics Meeting (ICHIM '01), Milan, Italy, Sep. 2001.

Eavesdropping on Electronic Guidebooks: Observing Learning Resources in Shared Listening Environments

Allison Woodruff, Paul M. Aoki, Rebecca E. Grinter, Amy Hurst, Margaret H. Szymanski, and James D. Thornton, Palo Alto Research Center, USA

Abstract

We describe an electronic guidebook, *Sotto Voce*, that enables visitors to share audio information by eavesdropping on each other's guidebook activity. We have conducted three studies of visitors using electronic guidebooks in a historic house: one study with open air audio played through speakers and two studies with eavesdropped audio. An analysis of visitor interaction in these studies suggests that eavesdropped audio provides more social and interactive learning resources than open air audio played through speakers.

Introduction

Previous research suggests that users of electronic guidebooks prefer *open air audio* delivered through speakers to audio delivered through a headset (see, e.g., Kirk, 2001; Woodruff, Aoki, Hurst, & Szymanski, 2001). The well-known visitor desire for social interaction (Hood, 1983) is a key reason for this preference: when visitors use open air audio, they can listen to content together and discuss it, whereas headsets often isolate visitors into experiential "bubbles" (Martin, 2000). However, open air audio is problematic when many visitors are present in the same location, as has been confirmed by informal experiments conducted by commercial audio guide vendors (L. Mann, Antenna Audio, personal communication).

We describe an alternative mechanism for sharing audio. This mechanism, which we call *eavesdropping*, preserves the social interaction enabled by open air audio while avoiding the audio "clutter" that open air audio necessarily entails. In our system, visitors independently select objects in their guidebooks and listen to the audio content through one-ear headsets; these headsets allow them to hear each other speak and interact conversationally. Further, wireless networking enables visitors to optionally listen to their companion's guidebook in addition to their own. The intimate, often directed, nature of the resulting shared audio context has led us to call the system *Sotto Voce*.

Our design is guided by the following principle: we want to support visitor interaction with three main entities that make demands on their attention. These entities are the information source, the visitor's companions, and the physical environment – "the guidebook, the friend, and the room" (Woodruff, Aoki et al., 2001). As we add capabilities that enhances visitor interaction with one entity, we must be careful that we do not compromise visitor interaction with the others (e.g., we do not want to improve visitor-visitor interaction at the expense of visitor-room interaction.)

To understand the impact of the eavesdropping mechanism on the overall visitor experience, we conducted two studies of visitors using the system to tour a historic house. We applied qualitative methods to the resulting data, including an analysis of visitor interviews and an applied conversation analytic study of recorded audiovisual observations. Because the eavesdropping was an optional feature that visitors could turn on or off at will, we observed several categories of use, e.g., pairs of visitors who did not use eavesdropping, pairs of visitors who used eavesdropping intermittently, and pairs who engaged in continuous mutual eavesdropping.

In this paper, we focus on the visitors who engaged in mutual eavesdropping, which is the category that

BEST COPY AVAILABLE

most closely approximates open air audio. We compare the typical behavior of these mutual eavesdroppers to that of visitors in a previous study who used open air audio to create a shared listening experience (Woodruff, Aoki et al., 2001; Woodruff, Szymanski, Aoki, & Hurst, 2001). (The three studies are summarized in Table 1.) Most of the discussion is based on analysis of the observational data. We observe that mutual eavesdroppers had a different activity structure and were more mobile than visitors who used open air audio. As a result of these changes, mutual eavesdroppers had increased resources for engaging in interactive learning: they had richer and more extensive social interaction, and they had more resources for physically exploring their environment. For example, visitors had more substantive discussion in response to guidebook descriptions, and they were more likely to discuss objects not described in the guidebook. Given the importance of social learning in the museum environment (Falk & Dierking, 2000), the preliminary evidence presented here is encouraging and suggests further avenues for work along these lines.

The remainder of the paper is organized as follows. First, we discuss the design of SottoVoce. Next, we describe the method employed in our user study. We then turn to findings. These are divided into the impact of the design on visitor behavior and the implications of these behavioral changes for visitors' learning resources. After discussing related work, we summarize our findings and describe future directions.

Prototype Design

In this section, we discuss the design and implementation of the guidebook device, key aspects of its user interface, the design goals for the audio environment, the eavesdropping mechanism, the audio delivery mechanism, and the construction of the audio content. The design is the same as that used in Study 2, reported in (Aoki et al., 2002), but we briefly discuss it here to provide context. Overall, visitors have a positive response to the guidebook and report that it is easy to use (Aoki et al., 2002; Woodruff, Aoki et al., 2001).

Guidebook device. We implemented the device using the Compaq iPAQ' 3650 handheld computer, which includes a color LCD touchscreen display. With an IEEE 802.11b wireless local-area network (WLAN) card, the device measures 163mm x 83mm x 34mm (6.4" x 3.3" x 1.3") and weighs 368g (13 oz.).

To support eavesdropping, paired devices communicate over the WLAN using internet protocols (UDP/IP). The audio content is the same on all devices, so the devices send and receive control messages ("start playing clip 10," "stop playing clip 8") rather than waveform audio. Since our goal is to enhance co-present interaction, the device does not support remote voice communication.

User interface. This part of the system is very similar to that used in previous studies, and its design rationale is more thoroughly described elsewhere (Woodruff, Aoki et al., 2001). Individual visitors obtain information about objects in their environment using a visual interface. This helps visitors maintain the flow of their visual task (looking at the room and its contents), which tends to reduce demands on user attention. The interface resembles a set of Web browser imagemaps; at a given time, the visitor sees a single photographic imagemap that depicts one wall of a room in the historic house (Figure 1, center). Visitors change the viewing perspective (i.e., display a different imagemap) by pressing a hardware button. When visitors tap on an imagemap target, the guidebook plays an audio clip that describes that object. Many, but not all, of the

Table 1. Summary of studies conducted.

Study	Audio sharing mechanism	Participants		Previous papers
		Recruited	Public	
1	Open air	14		(Woodruff, Aoki et al., 2001; Woodruff, Szymanski et al., 2001)
2	Eavesdropping	12		(Aoki et al., 2002)
3			47	--

Table 1

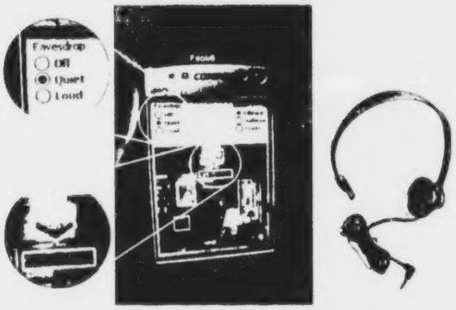

Figure 1. Electronic guidebook and headset.

objects visible on the screen are targets; to help visitors identify targets, the guidebook displays *tap tips* (Aoki, Hurst, & Woodruff, 2001) – transient target outlines that appear when the user taps and fails to "hit" a target (Figure 1, bottom left). A demonstration of the visual interface is available online (http://www.parc.com/guidebooks/).

Audio design goals. Results from Study 1 suggested several design criteria. Visitors want to be able to share audio descriptions and converse. At the same time, visitors want to retain personal control over the selection of descriptions. Further, the design needs to facilitate the ability of visitors to explore their physical environment, and the design needs to be sufficiently lightweight that it makes minimal demands on the users' attention. Finally, the design needs to be feasible in public environments with many visitors. These criteria ruled out a number of options like open air audio (which is not feasible for large numbers of visitors) or splitters that allow two visitors to listen to audio from a single device (which restrict visitor movement and do not allow visitors individual control over the audio content to which they are listening). The eavesdropping model described below is an alternative that meets all of the criteria.

Eavesdropping. In concrete terms, paired visitors share audio content as follows. When visitor A selects an object on her device, she always hears her own audio clip. If A is not currently playing an audio clip, but her companion B is, then B's audio clip can be heard on A's device. In other words, audio clips are never mixed, and A's device always plays a personal clip (selected by A) in preference to an eavesdropped clip (selected by B). Audio playback on the paired devices is synchronized; if A and B are both listening to their own clips and A's clip ends first, A will then hear the remainder of B's clip as if it had "started in the middle." To control a device's eavesdropping volume (i.e., the volume at which A hears B's clips), the interface includes three option buttons: "Off," "Quiet" and "Loud" (Figure 1, top left). "Loud" is the same as the volume for personal clips.

In abstract terms, eavesdropping provides a relatively simple *audio space* model (Mackay, 1999). We did consider other options, such as a telephony-like connection model in which visitors would independently initiate and terminate audio sharing sessions with their companions. We also considered email-like asynchronous models in which visitors would send and receive audio clips at their convenience. We rejected more complex abstractions that involved multiple actions (send/receive, connect/accept/reject, etc.) because we believed that the necessary interface gestures would distract visitors from their experience with the environment and their companions. In the audio space model, sharing requires no gestures of its own. To "receive," a visitor merely sets the eavesdropping volume. To "send," a visitor simply selects an object; playing a description has the side effect of sharing it, if the companion chooses to eavesdrop. The audio space model has the further advantage that it supports simultaneous listening, which enhances social interaction by creating the feeling that the content is part of a shared conversation (Woodruff, Szymanski et al., 2001).

Audio delivery. Visitors hear descriptions through headsets. We conducted a small study (n=8) to identify headsets that would allow visitors to converse and that visitors would readily accept (Grinter & Woodruff, 2002). Based on this study, we chose commercial single-ear telephone headsets, locally modified by the removal of the boom microphone (Figure 1, right). This configuration leaves one ear available to hear sounds from the external environment, and visitors find the over-the-head design desirable because it is familiar and gives them the sense that the headset is securely attached.

Audio content. The prototype contains descriptions of 51 objects in three rooms of the house. In most regards, the descriptions are recorded along principles described in (Woodruff, Szymanski et al., 2001). The audio clips vary in length between 5.5 and 27 seconds, with the exception of one story

30

that runs for 59 seconds. The clip length is much shorter than conventional audio tour clips, which often run to 180 seconds, and is intended to facilitate conversation by providing frequent opportunities for visitors to take a conversational turn.

Since we use single-ear headsets, both personal and eavesdropped audio content are necessarily presented in the same ear. We distinguish the two types of content using two mechanisms. First, we apply a small amount of reverberation to the eavesdropped audio. A single earphone cannot effectively deliver spatialized audio (Blauert, 1997), but can support other sound effects; we chose reverberation after conducting user tests (n=6) involving scenario-based tasks using the guidebook. Second, the default eavesdropping volume ("Quiet"), which is most frequently used by visitors, is softer than the personal volume.

Method

We have conducted three major user studies at Filoli, a Georgian Revival historic house located in Woodside, California (http://www.filoli.org/). Study 1 used an earlier version of Sotto Voce that supported open air audio, whereas Studies 2/3 used the current version of Sotto Voce that supports eavesdropping as described in the design section of this paper. Study 1 and Study 2 involved previously recruited participants on days the house was closed to the general public, whereas Study 3 involved 47 visitors recruited on-site on days the house was open to the general public. (Again, these studies are summarized in Table 1.)

Because the participants and procedures for Study 1 and Study 2 have been reported previously, below we report only the participants and procedure for Study 3. We then discuss our analytic methods, which were the same in all studies.

Participants. In Study 3, we observed 20 pairs, one group of three, and one group of four using the guidebooks. These pairs and larger groups were comprised of visitors who had come to Filoli together, e.g., mother/daughter or friend/friend pairs. The majority of visitors had not previously used a handheld device. The visitors covered a wide range of ages: the youngest visitors were in the "18-29" age range, and seven visitors who used the guidebook were "over 70." (While we had several children test SottoVoce in the first and second studies, visitors from the ages of approximately 5-17 are quite rare at Filoli unless they are visiting with a school group.)

Procedure. Visitors to the house were recruited at the entrance to the Library, the first room discussed in the guidebook. After signing consent forms, visitors were fitted with a wireless microphone, given guidebooks, and trained in their use. Next, they visited the three rooms for which the guidebook had content. When they finished using the guidebooks, they participated in a semi-structured interview.

The visitors' conversation and comments during the interview were recorded using the wireless microphones; the visitors were videotaped by fixed cameras while using the guidebooks (all visitors to the house were notified that videotaping was in progress); and the visitors' use of the guidebooks was logged by the device.

Visitors typically spent about 15 minutes using the electronic guidebooks. Their participation in the study took approximately 30-45 minutes; no time limits were imposed during any portion of the procedure.

Analysis. We analyzed the data in several ways. For example, we transcribed and analyzed the interview data to examine the visitors' attitudes and feelings about the technology and their experience. The majority of the findings presented in this paper are based on another method we used, *conversation analysis* (Sacks, 1984).

Conversation analysis is a sociological method used to examine naturally occurring social interaction to reveal organized patterns.

To find such patterns, conversation analysts study collections of interactive encounters and identify *sequences* of actions that were recurrently made by the participants. Actions in our context might include making a verbal utterance, pointing at an object, or selecting a description.

To this end, we create a composite video of visitors and their guidebook screens and audio (re-created from the guidebook activity logs). We then transcribe the actions taken by visitors, including dialogue, and look for recurring patterns to identify visitors' systematic practices.

From Open Air To Eavesdropping: Changes In Visitor Behavior

In this section, we compare the behavior of the pairs who chose to use mutual eavesdropping in Studies 2/3 to that of similarly engaged pairs who used open air audio in Study 1.

Specifically, we discuss the structure of the visitors' interactions and their physical mobility. The effect of these aspects can be identified in the visitors' learning-related behavior, which is the subject of the following section.

Changed Activity Structure

Visitor activity was structured very differently with eavesdropped audio than with open air audio. The new structure had a lower coordination cost, demanding less attention. The decreased attention burden was reflected in the visitors' interactions.

In all of the studies, a single overall structure pervaded the interactions. Specifically, they exhibited the sequential, multi-phase organization known as storytelling in the conversation analytic literature (Sacks, 1974); as part of this organization, visitors created a conversational role for the audio descriptions, i.e., they treated the guidebook like a "third party" taking an extended conversational turn (Aoki et al., 2002; Woodruff, Szymanski et al., 2001). Paired visitors entered a state of engagement at the beginning of a given storytelling sequence; levels of engagement generally rose and then fell over the course of a given sequence; and visitors then had the options of dis-engaging (resulting in independent activity), remaining engaged in shared activity, or maintaining a nascent engagement in expectation of subsequent re-engagement (Szymanski, 1999).

With open air audio, visitor interactions tended to focus on choosing individual objects and coordinating with their companions to listen to the descriptions. This setup, repeated for each sequence, focused more attention on coordination activity than seems necessary or desirable. However, the open air audio did afford the opportunity to participate in shared responses to the "story," motivating the visitors to begin setup for another sequence.

By contrast, participation in mutual eavesdropping created an ongoing assumption that the couple would continue in the shared activity. This supposition of continuing shared activity meant that setup tended to be cursory. Further, while open air audio was primarily conducive to follow-up discussions that related directly to descriptions, mutually eavesdropped audio was conducive to many diverse types of follow-up sequences such as discussion of objects not described in the guidebook.

The change in activity structure had at least two beneficial effects.

- First, by reducing the effort needed to choose and listen to descriptions, mutual eavesdropping freed visitors to direct more attention to meaningful interactions with their environment and their companions (i.e., away from the guidebook and routine coordination). In other words, the reduction in low-quality coordination talk meant that visitors had more time to investigate the room and its contents and that a higher proportion of talk tended to focus on topics of substance.

- Second, since the new activity structure supported more diverse types of sequences, visitors were more likely to pursue new topics or investigate objects not described in the guidebook.

Increased Mobility

Visitors in Studies 2/3 were noticeably more mobile during periods of engagement. In Study 1, the open air audio was played at a low volume, so any movement that changed the relative position of the visitors could cause significant sound attenuation due to distance or blockage (e.g., due to interposed obstacles – even changes in body orientation could cause the audio to be blocked). As a result, couples tended to remain close together and stationary while sharing audio descriptions. See Figure 2a, in which a grandmother is bending over to listen to the audio description that her granddaughter is playing of the portrait over the fireplace. Note how

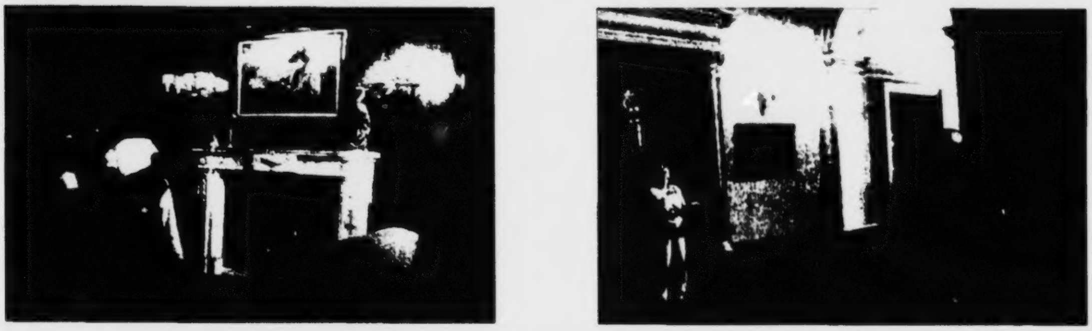

Figure 1. Comparison of visitor mobility patterns.

(a) Visitors standing close together when using open air audio.

(b) Visitors standing far apart when using mutually eavesdropped audio.

this position prevents her from examining the painting while she listens. In Studies 2/3, visitors were less constrained. Because movement could not attenuate the audio information, visitors could separate from each other physically while listening to descriptions and remaining engaged. See Figure 2b, in which both visitors are listening to a description of the marble staircase. While both visitors are examining the staircase, they have each chosen different vantage points. However, this positioning does not compromise their social connection: when the audio description reveals that only the first four steps are actually solid marble, the male visitor looks to his companion and she laughs, even though they are not standing together.

The increased mobility resulting from use of mutual eavesdropping took many forms. We observed several common behaviors that rarely, if ever, occurred with open air audio. For example, visitors would often walk together while a description was playing, e.g., to approach the object being described. In other cases, a single visitor would walk closer to the object currently being described while their companion remained stationary. In still other cases, a visitor would investigate a different object from the one currently being described and then rejoin the companion.

From Open Air To Eavesdropping: Increased Resources For Learning

The study observations provide evidence that both of the factors described in the previous section – the changed activity structure and increased mobility during engagement – improved the learning environment. Here, we discuss two learning-related resources that were enhanced by the guidebook: the nature of the visitors' social interaction and their opportunities for exploring the room and its contents.

Our analyses are based upon a collection of transcribed excerpts from which the following extracts have been derived. These extracts are meant to exemplify and highlight specific behaviors rather than to illustrate the organization of the visitors' interactions (space limitations preclude the use of representative excerpts of this kind).

Table 2 summarizes the notation used in this section. For clarity of presentation, the extracts have been simplified to use conventional capitalization and punctuation (e.g., commas and periods).

Note that the discussion in this section is limited to the *increased availability* of *learning resources*. A claim of *increased learning* would require a different type of study, e.g., one that measured the visitors' knowledge before and after their visit. Such a study is beyond the scope of this paper.

Table 2. Summary of transcription notation.

X: **X-PDA:**	Visitor X is speaking. ((comment or action)) *Visitor X's guidebook is speaking.*
°°whisper°°	Speech at reduced volume.
emphasis	Emphasis in speech.
(n)	A conversational pause of *n* seconds. (.) indicates a "micropause."

Depth and Length of Social Interaction

When using mutual eavesdropping, visitors responded more fully to audio descriptions. Visitors were also more likely to discuss features of the object not mentioned in the description or to discuss objects that were not described in the guidebook at all.

Each of these phenomena represents a way in which visitors collaboratively built on the shared audio descriptions, working together to construct mutual learning resources that broaden, deepen or expand their discussion of the room's contents. The importance of such social learning, particularly (but not limited to) conversation, has been widely supported in the visitor studies literature (see, e.g., (Falk & Dierking, 2000; Russell, 1994)). Social interaction around artifacts affords the "opportunity for the visitor to make connections with familiar concepts and objects" (Hein, 1995); adding resources for interaction adds more such opportunities. The remainder of this subsection gives some examples, linking the behavioral changes of the previous section to the construction of learning resources.

Characterizations. With mutual eavesdropping, response to an audio description was likely to be more reflective and include a physical focus relative to the object being described. With open air audio, visitors would often have a very minimal response, and even the more substantive responses were generally limited to reactions to the description that had just occurred.

The following extracts are representative of this effect. Consider V and W (Extract I, Study 1) who were listening with open air audio.

> **V-PDA:** Many of the top shelves contain false books. They are lighter than normal books, so they reduce the stress on the bookcases. Many are made of greeting cards, clothing, fabric, et cetera.
> **W:** Eh hah, that's a riot. ((W looks at V and smiles)) (0.2) They're just for looks.

Extract I.

While V and W do share a response, the substance is limited to a single paraphrase of the audio description, analogous to "text echo" of exhibit labels (McManus, 1989) (though possibly more affective, because of the audio delivery).

By contrast, consider an interaction in which J and L (Extract II, Study 2) were mutually eavesdropping.

> **L-PDA:** . . . All of the architectural features of this room, including the walnut panelling, are modelled on an 18th century British library. In the original library, each of the outlined panels would have contained framed pictures.
> **L:** Really.
> **J:** Yeah.
> **L:** That's a lot of pictures. ((points at wall and sweeps arm across walls))
> **J:** That's a lot of pictures. ((nods "yes"))(.)
> J: That would've been very cluttered.

Extract II.

This interaction is richer in many respects than that shown in the previous extract. J and L both speak, taking more turns to discuss the description than V and W. By making a statement about the number of pictures, L reinforces for J a quantitative observation that is not made in the description itself. By gesturing at the many empty wall panels, L adds physical, spatial, and visual elements to the experience, linking both visitors to a vision of the "original" library that overlays their actual surroundings. J agrees with the quantitative statement, pauses, and then responds by saying it would have been "cluttered," a qualitative assessment that indicates that she has in fact visualized the room as it might have been.

This increased reflection on descriptions was evidenced in many ways. Visitors worked together more to understand descriptions, e.g., a visitor would sometimes express confusion about a description and the companion would help explain it. Further, with mutual eavesdropping, visitors more frequently branched off into sequences that were not directly related to the description of the content. For example, they might point out a specific physical feature of the object that was not mentioned in the description, or discuss some way that it related to their own life.

Additionally, visitors showed more evidence of establishing complex relationships between objects. One visitor pointed at a series of paintings on different walls, saying, "Okay, so that's his wife, and that's his mother, right?" Or consider T's comments (Extract III, from Study 2) when he first enters a particular room. His statements indicate that he has constructed a category of "secret cabinets" that occur in this house and that he is alert to instances of this category as he moves from room to room.

T: Ah, more secret cabinets.
(0.4)
T: I like that a lot about this house. ((walks into the bar closet))

Extract III.

Interactions displaying this kind of orientation – i.e., at the granularity of a thematic collection rather than a single object – almost never occurred with open air audio.

Moreover, mutually eavesdropping visitors often discussed objects that were not described in the guidebook, unlike open air audio visitors. The following sequence, in which J teaches L about a plant, occurred immediately after they finished their response to a description:

J: Okay, your- your test for the day, what's that one? ((points to plant))
(0.2)
J: The plant.
(0.4) ((L leans in to look))
L: Morning glory. Eh heh heh heh, I don't know, what is it?
J: I think it's a mandevilla vine, but I'm not sure.
L: God, I can't believe you know that.

Extract IV.

Reasons. Both of the behavioral changes resulting from use of mutually eavesdropped audio had impact on social interaction. The primary factor was the new activity structure, which allowed more space for reflection and for visitors to initiate new conversational sequences that were not structured around the audio descriptions. Increased mobility constituted a secondary factor. Visitors would often start descriptions while they were far away from objects. As mentioned above, visitors were unlikely to walk toward the object while the description was playing with open air audio. However, with eavesdropped audio, they were more likely to approach the object; being close to the object when the description ended gave them more opportunities to observe and discuss its specific features.

Expanded Resources for Physical Exploration

With mutually eavesdropped audio, the examination of objects was more frequently occasioned by their presence in the *room* rather than their presence in the *guidebook*. Once visitors began to examine an object, they might discuss it or play a description of it if one were available.

This implicit shift in emphasis from the guidebook to the room as the impetus for exploration is important because it shifts the visitor's role. It is broadly (though perhaps not universally) accepted that learning is enhanced by enabling visitors to navigate the museum without leading them through it (Falk & Dierking, 2000). However, even "free choice" navigation can be constrained by, e.g., which objects have descriptive content associated with them. Visitor behavior indicates that use of mutual eavesdropping increased the guidebook's utility as a reference (an adjunct to the room) as opposed to an inventory (a directed guide to the room).

Characterizations. In the study using open air audio, examination of objects often began with objects contained in the guidebook and proceeded by spatial locality. That is, visitors tended to switch the visual interface to a given wall and then look at the objects in the guidebook that interested them on that wall. Object choice was often based on targets seen in the visual interface or on short-term memory of such targets.

In the eavesdropping studies, the next object to examine was less frequently chosen based on availability in the guidebook. (In many of these cases, we know that the examination was prompted by the room rather than the guidebook because the objects were not described in the guidebook. In the other cases, the visitors spoke their thoughts aloud – which was entirely self-prompted since none of the studies involved a speak-aloud protocol.) Instead, visitors would encounter objects in their field of view, e.g., objects that were near an object they had just examined, or they would deliberately examine sequences of objects they perceived as being related.

For example, in Extract III, T walks into a new room, notices the bar closet and actually walks into it. *After* this, his companion D finds the description in the guidebook and plays it. Note that because the sound does not attenuate, the visitors can listen to the description together while T stands inside the tiny closet and D stands outside.

Reasons. While the same resources were available with open air audio, they were used much more frequently in the mutual eavesdropped case due to the changed activity structure and the increased mobility in the room. Specifically, the mutually eavesdropped audio was more conducive to sequences that were not directly responsive to guidebook content; visitors were generally more open to external triggers with the new activity structure. Visitors acted in a manner more consistent with "Let's see what's here in the room" than with "Let's see what's here in the guidebook." Further, visitors had more attention to give to the room due to the reduced attentional demands and wandered more in the room due to increased mobility, so they were more likely to encounter and investigate objects.

Related Work

Our work draws together three main areas of research. Space limitations preclude an extended discussion; additional references are contained in (Aoki et al., 2002; Woodruff, Aoki et al., 2001; Woodruff, Szymanski et al., 2001).

Interaction in museum settings. The importance of social interaction to museum visitors is well known (e.g., (Hood, 1983)). There are two types of studies of particular interest. McManus observed visitor usage of text labels; she noted that visitors were inclined to treat exhibit labels *as conversation* to which they had been party (McManus, 1989). A number of studies of museum visitors have been conducted using methods derived from conversation analysis (see, e.g., (Falk & Dierking, 2000), Ch. 6, and (vom Lehn, Heath, & Hindmarsh, 2001). These studies focus on talk, interaction and learning in conventional environments; here, we have focused on the effects of electronic guidebooks on social interaction and learning resources.

Electronic guidebooks. The cultural heritage community has formally studied electronic guidebooks for many years (Screven, 1975). Related work in HCI has focused on aspects such as location-aware computing (Abowd et al., 1997), and only recently have significant user studies been reported (e.g., (Cheverst, Davies, Mitchell, Friday, & Efstratiou, 2000)). The HCI studies focus on system design and evaluation; here, we focus on the effects of our system on visitor interaction.

Media and interaction. There is an extremely rich literature on collaborative multimedia environments; of particular interest are media spaces (Mackay, 1999). Many of these systems have been evaluated, but most apply either ethnographic techniques or quantitative methods to studies of installed workplace systems. In this study, we apply conversation analytic techniques to the study of a mobile, leisure-activity system that provides shared access to application content.

Conclusions

In this paper, we have described an eavesdropping mechanism that allows visitors to listen to each other's guidebooks. Our findings show that mutual use of this eavesdropping mechanism can lead to increased learning resources as compared with the use of speakers in open air: couples using mutual eavesdropping in Studies 2/3 had more substantive interactions and exhibited an increased awareness of the room and its contents when compared to those using open air audio in Study 1.

New work is addressing some of the open issues from this study. We are preparing a discussion of the ways in which the visitors creatively used our eavesdropping mechanism for tasks other than enhancing their social interaction, e.g., for monitoring their children. We are also planning an experiment using bone conduction headsets that can provide binaural audio without occluding the ears.

Acknowledgements

We are deeply grateful to Tom Rogers and Anne Taylor of Filoli Center for their assistance with this project. Amy Hurst performed this work during an internship from the College of Computing, Georgia Institute of Technology.

References

Abowd, G. D., Atkeson, C. G., Hong, J., Long, S., Kooper, R., & Pinkerton, M. (1997). Cyberguide: A Mobile Context-Aware Tour Guide. Wireless Networks, 3(5), 421-433.

Aoki, P. M., Grinter, R. E., Hurst, A., Szymanski, M. H., Thornton, J. D., & Woodruff, A. (2002). Sotto Voce: Exploring the Interaction Between Conversation and Mobile Audio Spaces. In Proc. ACM SIGCHI Conference on Human Factors in Computing Systems. New York: ACM Press, to appear.

Aoki, P. M., Hurst, A., & Woodruff, A. (2001). Tap Tips: Lightweight Discovery of Touchscreen Targets. In Extended Abstracts, ACM SIGCHI Conference on Human Factors in Computing Systems. New York: ACM Press, 237-238.

Blauert, J. (1997). Spatial Hearing (rev. ed.). Cambridge, MA: MIT Press.

Cheverst, K., Davies, N., Mitchell, K., Friday, A., & Efstratiou, C. (2000). Developing a Context-Aware Electronic Tourist Guide: Some Issues and Experiences. In Proc. ACM SIGCHI Conference on Human Factors in Computing Systems. New York: ACM Press, 17-24.

Falk, J. H., & Dierking, L. D. (2000). Learning From Museums. Walnut Creek, CA: Altamira Press.

Grinter, R. E., & Woodruff, A. (2002). Ears and Hair: What Headsets Will People Wear? In Extended Abstracts, ACM SIGCHI Conf. on Human Factors in Computing Systems. New York: ACM Press, to appear.

Hein, G. E. (1995). The Constructivist Museum. Journal of Education in Museums, 16, 21-23.

Hood, M. G. (1983). Staying Away: Why People Choose Not to Visit Museums. Museum News, 61(4), 50-57.

Kirk, J. (2001). MUSEpad. Presentation at the Electronic Guidebook Forum, San Francisco, CA, October 2001.

Mackay, W. E. (1999). Media Spaces: Environments for Multimedia Interaction. In M. Beaudouin-Lafon (Ed.), Computer-Supported Cooperative Work. Chichester, UK: Wiley & Sons, 55-82.

Martin, D. (2000). Audio Guides. Museum Practice, 5(1), 71-81.

McManus, P. M. (1989). Oh, Yes They Do! How Visitors Read Labels and Interact with Exhibit Text. Curator, 32(3), 174-189.

Russell, T. (1994). The Enquiring Visitor: Usable Learning Theory for Museum Contexts. Journal of Education in Museums, 15, 19-21.

Sacks, H. (1974). An Analysis of the Course of a Joke's Telling in Conversation. In R. Bauman & J. Sherzer (Eds.), Explorations in the Ethnography of Speaking. Cambridge: Cambridge Univ. Press, 337-353.

Sacks, H. (1984). Notes on Methodology. In J. M. Atkinson & J. Heritage (Eds.), Structures of Social Action. Cambridge: Cambridge Univ. Press, 21-27.

Screven, C. G. (1975). The Effectiveness of Guidance Devices on Visitor Learning. Curator, 18(3), 219-243.

Szymanski, M. H. (1999). Re-engaging and Disengaging Talk in Activity. Language in Society, 28(1), 1-23.

vom Lehn, D., Heath, C., & Hindmarsh, J. (2001). Exhibiting Interaction: Conduct and Collaboration in Museums and Galleries. Symbolic Interaction, 24(2), 189-216.

Woodruff, A., Aoki, P. M., Hurst, A., & Szymanski, M. H. (2001). Electronic Guidebooks and Visitor Attention. In D. Bearman & J. Trant (Eds.), Cultural Heritage Informatics 2001 (Proc. 6th ICHIM). Philadelphia: A&MI, 437-454.

Woodruff, A., Szymanski, M. H., Aoki, P. M., & Hurst, A. (2001). The Conversational Role of Electronic Guidebooks. In G. D. Abowd, S. Shafer & B. Brumitt (Eds.), Ubiquitous Computing (Proc. 3rd International Conference). Berlin: Springer-Verlag, 187-208.

Can You See Me? Exploring Co-Visiting Between Physical and Virtual Visitors

Areti Galani and Matthew Chalmers, University of Glasgow, Department of Computing Science, UK

http://www.dcs.gla.ac.uk/equator/city.html

Abstract

We explore issues of social context and interaction between digital and physical museum visitors, using as a focus of discussion the *City* project, itself set within a larger interdisciplinary project called Equator. We look at collaborative environments that span different media, in particular handheld mobile devices, Web-based hypermedia and 3D virtual environments. We discuss two main strands in our research: the methods and results of two pilot visitor studies in two cultural institutions in Glasgow—the Lighthouse and the House for an Art Lover—and the development of our prototype system which establishes three-sided collaboration between physical, Web and virtual environment visitors. We then present preliminary results and issues arising from our on-going system development and user trials. We conclude with future plans for further system evaluation and deployment.

Keywords: Visitor studies, social context, virtual reality, wearable computers, hypermedia.

Introduction

In their definition of the interactive museum experience, Falk and Dierking (1992) described three key elements that influence the way visitors experience museums: the physical context, the personal context and the social context. Physical context mainly covers the physical layout of the space and has been extensively studied by designers (Communications Design Team, 1976) and more recently by space syntax theoreticians (Psarra, Grajewski & O'Neil, 2002). Personal context covers the prior knowledge of the visitors, their personal aims and expectations, and their current state of mind. Personal context and its influence on the learning experience have also been studied by means of evaluation teams and learning theories such as constructivism and Gardner's multiple intelligences theory. Last, but not least, social context covers the social interaction during the museum visit between the visitors and their immediate companions, as well as other visitors and museum staff. Several aspects of the social context in a physical museum setting have been examined, focusing mainly on school groups and families (Diamond, 1986; McManus, 1987; Falk & Dierking, 1992). More recently, social scientists in the *SHAPE* project also examined the social cues in the interaction with displays (Vom Lehn, Heath & Hindmarsh, 2001).

This paper further examines social context in museums, and its technological support. We are especially interested in social context among new audiences who visit museums via the Web and 3D virtual reality applications. While studies of traditional museum audiences have shown the importance of social context, new media as used in museums do not often support social interaction. In studying both traditional museums and new technologies, we aim to address issues such as how to integrate visits to the virtual museum with visits to the 'physical' museum, especially when a visitor may visit both. This is partly a response to the way that the number of digital visitors is steadily growing and, in some cases, outstripping the number of visitors to the corresponding physical museums (Lord, 1999). As a result, many wish to find ways to encourage the geographically distant digital visitor to become a physical visitor, and to encourage physical visitors to maintain a relationship with the museum after they walk out its door.

Our project, City, investigates ways to support and enrich context in museums, cultural institutions and the city—in particular, social context. Set within a research consortium called Equator (www.equator.ac.uk), our work is on context that involves something richer and more complex than a collection of isolated media and disjointed pieces of information. Computer scientists as well as museum professionals—for different reasons—tend to focus on the obvious differences between traditional and digital media, and treat each one independently. Here, a broader viewpoint takes account of their similarities and interdependencies.

We initially discuss our approach to understanding and working with social context, pointing out prior work in both traditional and digital media. We then present two parallel lines of work that share a theme: the architect, designer and artist, Charles Rennie Mackintosh (1868-1928). We report on studies of social interaction in existing museums and exhibitions, and on a system infrastructure and prototype to bridge digital and physical visits to an exhibition room and its collection of artifacts. This prototype supports interaction between people visiting or exploring the room, even though they use quite different media: wearable computers, hypermedia and virtual environments. Thus a visitor using any one of these styles can interact with other visitors using other media. We offer some initial observations on the use of our prototype system before outlining our ongoing and planned work.

Social context

We approach the issue of social context by considering interaction between visitors. We initially use a deliberately naïve categorization of context, dividing it in terms of time, space and medium. First, interaction may be synchronous or asynchronous. Secondly, it may occur locally, within the museum, or remotely, with at least one visitor being physically located beyond its walls. Thirdly, it may involve digital media such as Web sites and virtual environments, or traditional 'physical' media such as the museum building itself, displays and, of course, the artefacts of the museum collection.

One experiences synchronous social awareness when visiting a museum with some friends or as a member of a larger group: one's *co–visitors*. The design of the exhibition space, through its layout, displays, artefacts and supporting materials, clearly influences the visitor's movement and activity, but the way co–visitors move around the space is also influential. Other visitors who happen to be in the museum at the same time may also have an effect, as may those who visited at some time in the past.

This latter influence, asynchronous social awareness, is a key feature of the museum experience. Traditionally afforded comment books, visitors can write about what they did or felt during their museum experience. Any later visitor can read the comments; i.e. the communication is not generally personalized or directed towards any one particular reader. Recently entire exhibitions and displays have revolved around this notion of social context. For example, in London, visitors to the Science Museum (www.sciencemuseum.org.uk) can read other visitors' opinions on controversial issues, and in the interactive games of the Wellcome Wing can compare their results with those of previous visitors. Visitors also have the opportunity to create a Web page with the highlights of their visit, and access it later online with their friends and family. This raises the issue of whether and how to structure, curate or select from such contributed 'collections' as they grow in size and value.

Web sites may offer various forms of social awareness. These may be synchronous, as in chat rooms, or asynchronous, as in mailing lists and recommendations. Awareness much like that of comment books is afforded by, for example, the Amazon.com bookstore. One can leave comments about specific books, CDs and so forth, and can read comments from earlier visitors. However, dynamic and personalized information can be cheaply offered when compared to traditional media. For example, each Amazon visitor's profile of movement around and purchases from the site is dynamically combined with others in creating recommendations for the individual, i.e. selections from the many items sold by Amazon, based on the match between a visitor's profile and those of earlier visitors. Each movement and purchase leads to different personalized recommendations. The *Hippie* (Oppermann & Specht, 1999) and the *SottoVoce* projects (Aoki & Woodruff, 2000) have already explored many of the issues involved in the personalized delivery of content in a museum environment but, as far as we know, recommender systems have not yet been applied in museums.

Synchronous social interaction is also a key feature in virtual environments. The vast computer game industry is centred on such technology, and is now exploring the possibilities of collaborative game playing on-line. Furthermore, chat channels such as Active Worlds (www.activeworlds.com) combine directed textual communication with graphical representation of users as avatars. Museums and related institutions have explored the same medium, but for purposes such as the presentation of artefacts too fragile or numerous to be put on display and of reconstructed archaeological sites, and in supporting geographically distant visitors. Such a medium supports social interaction between visitors, as in the Virtual Leonardo project in the National Museum of Science and Technology in Milan (www.museoscienza.org), and the Van Gogh Museum in Amsterdam (www.vangoghmuseum.nl). Web users share a common representation of the museum and are afforded basic resources for social interaction: mutual visibility and audibility. Immersive VR technology, such as head–mounted displays and room–sized (and room–shaped) projection surfaces, is used in a number of cultural institutions, e.g. the Foundation of the Hellenic World (www.fhw.gr). Relatively little support is given in such systems for asynchronous awareness, however.

Most research on museums and the Web has, we suggest, tended to treat local and remote visitors in isolation from each other, and to treat traditional and digital media similarly. To our knowledge, there has been little or no work that bridges between local and remote, and between traditional and digital media. However, those who visit the digital museum may visit the traditional museum, and vice versa. Previously seen digital information may influence a visitor's interpretation of traditional displays and artefacts, and vice versa. New technologies let visitors to the traditional museum interact with digital visitors, and combine traditional and digital media, in the same experience.

In order to deepen our understanding of social context, we have undertaken a pair of ongoing studies of social awareness and context in traditional museums and cultural institutions. These are described in the next section. We then describe a parallel stream of activity within our project, building a technological infrastructure for co–visiting and synchronous social context that crosses or blurs the boundaries between visitors who are local and remote, and between digital and traditional media. As described in a later section, ongoing and future work is directed towards applying the results of our studies to our systems, especially with regard to more explicit support for the roles visitors take with regard to each other, and towards asynchronous interaction.

Studies

In collaboration with two institutions in Glasgow, The Lighthouse (www.thelighthouse.co.uk) and the House for an Art Lover (www.houseforanartlover.co.uk), we are carrying out two sets of visitor studies of qualitative character. Our methodology is influenced by recent qualitative museum visitor studies as well as by ethnographic and anthropological methods used in social sciences and recently in computing science. Our aim is to better understand the relationship between the visitor and the social context and environment of the visit. These studies, unlike the majority of the visitor studies conducted in the museum sector, do not look at interactions around specific displays or temporary exhibitions. Instead, both studies are situated in permanent exhibitions.

We believe that qualitative methods, such as naturalistic ethnography, in the exploration of visitors' interaction are necessary to achieve a better understanding of 'how visitors see things, and what meanings they give to their experiences, rather than simply to enumerate frequencies for pre-formed categories' (MacDonald, 1993). Furthermore, our corpus of observations can be used to address either of two related issues: the generation and delivery of content, and social interaction during the visit. The latter will be the main focus of the following discussion.

The two institutions were chosen because they share the same topic, Mackintosh, but they explore different ways of engaging the visitor. The Lighthouse has developed an interpretation centre, often simply called the Mack Room, with a number of original objects intermixed with 23 workstations and displays that convey a substantial amount of digital information. The House for an Art Lover is a recently constructed house, but built, decorated and furnished according to Mackintosh's entry to a 1901

40

Figure 1. An interactive display in the Mack Room is operated by a visitor with more background knowledge of the exhibition than his co-visitor.

design exhibition. It is widely perceived as a historic house attraction, and a visitor is offered a leaflet and an audioguide that describe the construction project and Mackintosh's work.

In the Lighthouse we watched visitors and in eight cases manually recorded their movements on a map. In the House for an Art Lover we conducted participant observation of visitors for six days within a period of two months. In both locations, we used photographs and notes to add to our record of visitor activity. We also obtained some other sources of information regarding visitor experience. We were offered copies of the Mack Room designer's architectural drawings, showing the expected flow of archetypal visitors around the exhibits in the Mack Room, along with notes on the expected experiences gained by following such paths. In the House for an Art Lover, we were given access to the results of a marketing survey of visitors.

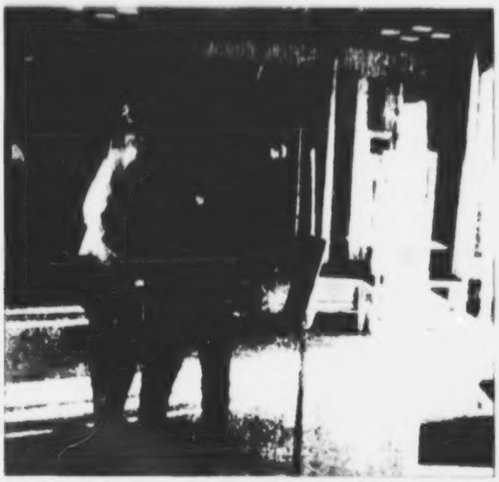

Figure 2. The visitor on the right, who has been in the House for an Art Lover several times before, points out and describes to her co-visitor a less obvious feature of the ceiling lights' design

In our studies we concentrated on 'casual visitors' (Falk & Dierking, 1992) in groups or as singletons, rather than the more structured visits of school groups and families. We focused on interpersonal interaction, and also observed how the information delivery media available in the gallery, such as audioguides, touch-screens, videos, labels and so forth, influenced this interaction.

We categorized the interaction between co-visitors into three main styles. In the first, co-visitors are tightly connected, staying together during their visit, and interacting with the same display at the same time. In the second style, co-visitors are loosely connected: they interact with different displays, but in the same area and thus stay relatively close to each other. Thirdly, co-visitors are independent navigators, following their own individual routes for the main part of the visit and meeting with each other only occasionally.

In tightly connected interaction, co-visitors stay close to each other as they actively and collaboratively interact with the displays, discussing artefacts and their descriptions. An important influence on this style of interaction is the background

Figure 3. Although close by each other, these two co–visitors in the House for an Art Lover are attending to different sources of information: an audioguide and a leaflet respectively

Figure 4. This couple is walking round the Mack Room together, but the woman on the right interacts with a touch screen while her companion, on the left, looks over her shoulder at a video display.

knowledge of visitors and their history of previous visits to the gallery. The most experienced usually leads the interaction by operating the touch screens (Fig. 1), or by pointing out details not obvious from the labels/commentaries (Fig. 2). Different co–visitors may undertake the leading role at different times during a visit.

The pace of the visit is an essential factor in this style. People usually share their interests in the displays and also share a pace of reading the available material. They also decide together where to go next. This style of interaction occurs quite often in galleries where information is not easily available or accessed, so people tend to move around together and help each other to interpret the exhibits.

In the second style of interaction, co–visitors are loosely connected. These visitors do not consistently interact with the same objects and displays, but stay close enough to maintain an awareness of each other. They usually exchange brief comments about what they have seen and point things out to each other. They may be very close to each other but engaged in quite different information, as in the leaflet reader and audioguide listener in Figure 3, and may attend to different displays at any given moment, as in Figure 4.

Visitors using audioguides usually experience this style of interaction. The audioguide often works as a resource for discussion despite the fact the device itself is not designed to support explicit interaction between visitors. This style of co-visiting also seems to favour gestural communication: a person may want to share an experience but not to interrupt the companion's experience, so favours hand gestures or eyebrow movements rather than verbal communication.

Independent navigation occurs in co–visitors who visit displays separately from each other. Each person has an individual pace, resulting sometimes in people waiting for the co–visitors outside the galleries, or 'slow' visitors cutting their visits short in order to catch up with others. Repeat visitors to the gallery often follow this style, and in some cases they leave the immediate space of the gallery and proceed to another room. Repeat visitors also tend to not take the audioguide at all or to consult it very infrequently compared to less experienced co–visitors.

This style of co-visiting includes many meeting points between the members of the group: sometimes accidental, sometimes deliberate. For example, two co–visitors may meet accidentally when their independent paths cross in a display that is of interest to both. In order to deliberately meet, one visitor may change navigation in the space, moving very quickly between points with good visibility in the gallery in order to find the co–visitor. Regardless of the way people meet, they discuss what they have seen in their individual tours, or suggest displays to each other—in some cases almost dragging their companions to show them a specific display. They also discuss the logistics of the visit; for example, where to go next, when they will have lunch, and so forth.

It is important to point out visitors do not rigidly conform to the same archetypal style of interaction during a visit. In all observed cases, co–visitors experienced their immediate social context through direct and close interaction at some times, and peripheral awareness at other times. Each co-visitor takes advantage of a dynamically changing set of resources. We identified mutual audibility, mutual visibility and shared content as essential resources in co-visiting. These three do not influence each of the interaction styles described above to the same extent.

Closely connected co–visitors use all three of the resources extensively. They have concrete knowledge of their companions' locations and orientations, which are often similar to their own. They talk almost constantly with each other, and share content whether in the form of touch screens, video screens or labels. They can share and explore interpretations of content seen or heard previously. Loosely connected co–visitors are also aware of each other's location and possibly orientation at any given moment. They use verbal or gestural communication, but they do not necessarily share content synchronously. Finally, independent navigators are only peripherally aware of their companions' locations. Such a visitor will usually inform co–visitors in a room when about to leave it, and indicate a destination, if their mutual awareness extends beyond a single room. They do not constantly share other resources but are potential users of all of them during meetings.

The Current System

In parallel with our studies, we have developed a system to explore social context that bridges or blurs the boundaries between visitors who are local and remote, and between digital and physical. The system involves three types of visitors: someone walking around the Mackintosh Room with a wearable computer, someone remote from the Mack Room but moving through Web pages related to the room, and another remote visitor using a VR model of the Mack Room.

A set of scenarios (Chalmers, 2001; Galani, 2001) was used to further explore the concept of co-visiting in museums and the city, and as a focus for system design. In our current scenario a visitor to Glasgow, Vee, visits the Mack Room in The Lighthouse. While browsing the displays, she invites her friend Anna, who is some distance away, to join her on the visit. Anna joins in by accessing an on-line 3D model of the room. A third friend, Dub, is then invited; he joins them by accessing the Web site for the room. The scenario introduces three different visitors who are remote from each other, and who have access to different technologies. Such scenarios have proved an effective way for a diverse team of researchers to communicate with each other, allowing us to shift our focus between studies, design, implementation and evaluation while maintaining a common context.

Vee, or the visitor in the role of Vee, has a handheld computer, an HP Jornada, which supports wireless network communications. She also has a jacket—Bristol University's *CyberJacket*, shown in Figure 5—that has within or on it several devices: an ultrasound detector to register position within the room

(Randell & Muller, 2001), an electronic compass to show orientation, and a battery. Cables within the jacket connect these parts to each other and to the handheld. The recorded position is updated every few seconds and sent off to another 'server' computer (hidden behind the reception desk of the Mack Room). The handheld computer can be used to browse pages about the Mack Room's displays and artefacts based on the exhibition's catalogue and delivered from the Web server on the local network. A number of those pages generally describe thematic zones within the room. We support automatic triggering of those pages: if Vee walks into a thematic zone within the room, a Web page for that zone is automatically displayed.

Dub has a normal Web browser running on a laptop computer. We currently set this laptop in The Lighthouse but away from the Mack Room, and connect it to the local network. He can browse pages about the Mack Room, and by tracking the zone associated with each page displayed, we imply a 'position' in the room for him.

Anna uses a VR model of the Mack Room as well as a Web browser. Unlike the Web visitor, Anna's position is her position within the virtual environment. We initially experimented with the room-sized immersive VR display (ReaCTOR) at University College London (UCL), but current experimentation happens within The Lighthouse, where we use another networked laptop showing simpler 'desktop VR' graphics.

Each visitor has a microphone and headphones, and can talk with the other two. We initially used mobile telephones to support conversation, but now have audio connections within our system. Each visitor can see personal location as well as locations of the others. Vee and Dub each have a map showing the Mack Room set within their Web page and showing visitors' ongoing positions as dots, as in Figure 6. Anna has her co–visitors shown as avatars, as in Figure 7.

A server computer assists the visitors' personal systems. This runs a Web server based on the *Linky* software (www.equator.ecs.soton.ac.uk/technology/leaky/leaky.shtml) developed by the University of Southampton to deliver hypermedia tailored to the context of the visitor. At the moment, the 'context' used by this software is the visitor's location and the size of the display. Linky models associations between artefacts, zones and hypermedia content and, when delivering a Web page for such a location, it can use these associations to add descriptions of other things to see and places to move to. The server computer also runs software, called *EQUIP* (www.equator.ac.uk/technology/equip/index.htm), that collects the positions of the various visitors, relates them to zones, and triggers any automatic display of Web pages, maps and avatars. EQUIP also handles audio streams between visitors.

Figure 5. Bristol University's CyberJacket is worn by the 'physical' visitor in the Mack Room (on the right of the picture), and has location and orientation sensors for a handheld computer.

Experimenting with the System

Our approach to system development is to carry out short tests or trials, where minor system refinements can be made, in between each major developmental step. The first of these was carried out in two rooms in UCL when we had established an initial but rough communication between the wearable system and the VR system, i.e. between Vee and Anna. We have just begun a second trial in the Lighthouse, with the system supporting three–way communication between Vee, Anna and Dub.

In the UCL trial, we invited colleagues external to the City project to take on the roles of Vee and Anna, so that we might gain initial observations of

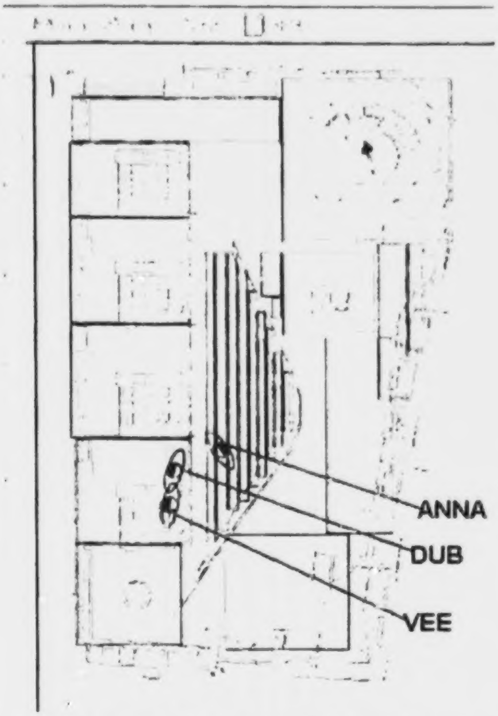

Figure 6. ***The web visitor, Dub, sees a map of the Mack Room showing the positions of himself and his co–visitors, and the outlines of thematic zones.***

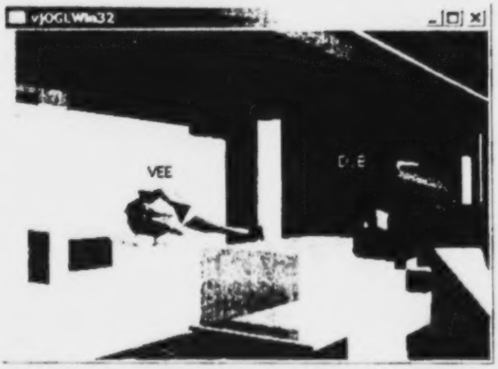

Figure 7. The view of the VR visitor, Anna, shows the Mack Room and avatars of her co–visitors.

interaction between a visitor in a virtual environment and a visitor with a handheld device in a physical environment. We videotaped the room where Vee was, and the VR as seen by Anna. We used a smaller version of the Mack Room and a number of objects related to Mackintosh, such as posters and postcards. This trial took place in UCL, rather than the Lighthouse, to let us experiment with an immersive VR display and to use UCL's infrastructure for recording activity in a virtual environment. One of the participants also wrote a lengthy feedback document after the event (Stenton, 2001).

In general, the participants felt that being in different geographical spaces did not make mutual awareness problematic. They extensively used the audio connection during the trial, especially when there were issues of shared reference to resolve. With limited facilities for sharing body orientation, they often spoke to each other to confirm or question exactly what they were facing and what they were referring to. Participants' engagement in the shared experience was shown by their use of vocabulary and body movements used in their everyday physical interactions, and by their awareness of each other's relative positions. Two instances of the latter involved the participant Phil (as Vee) who felt that he was in the way of Bill (Anna), and stepped back to free the view:

> P: Can you read this poster on the wall?
>
> B: Where? Which one are you looking at? That one...
>
> P: Oh, I'm sorry, I'm in the way mate. Can you read it now? Can you read the Willow Tearooms?
>
> [...]
>
> B: ...interferes with the wall. I believe I am too close to the physical wall.
>
> P: Oh yeah – I'll go out of the way. There you go. Don't go through the wall.
>
> B:That's the Willow Tearooms again.
>
> P: Ah, the Willow Tearooms... Sorry, you must be standing on my foot then
>
> B: I believe I'm sitting inside you. Yes. Its' rather...
>
> P: Oh... inside me – what a thought!

The video footage revealed that in the first of these cases, Phil was not quite in Bill's way. However, Phil's concern and motion to avoid blocking Bill's view indicated a sense of shared experience and presence. As Phil later wrote:

> Clearly our dialogue was one of discovering the shared experience rather than sharing the visit. A defining moment for me was the request from my virtual partner, "Get out of the way; I can't see the poster." At this point it felt like we were in the same room rather than sharing information about two identical rooms in different places (Stenton).

With regard to the participants' use of the handheld and the VR equipment respectively, we observed a separate set of issues. One of the common visitor gestures in a traditional museum setting is to point at something. In the case of the VR, this is tracked and displayed; i.e. the avatar has a head and a hand. We observed, however, that when using the wearable and handheld, the visitor used his free hand to point at things. This gesture went unnoticed by the system and consequently by his companion. On the other hand, the digital visitor had great difficulty reading textual information due to limited resolution in the virtual environment. Barbieri and Paolini (2000) examined the same problem with regard to desktop VR in the Virtual Leonardo project, and they used 2D graphics in the delivery of textual information. We use the same solution in the desktop VR version of our system, although we are also considering the use of a second handheld device to display textual information to the VR visitor, Anna.

Field trials inside The Lighthouse are, at the time of writing, still in progress. The next section outlines this and other ongoing work.

Ongoing and Future work

So far we have discussed two parallel strands of our work. The studies and system development have raised our awareness of social context issues, and of the possibilities and the limits of our technology. While the studies continue, the system work has subdivided into two strands. First, our technological infrastructure is being refined and expanded to support future systems where we support visitors moving between buildings and exhibitions across the city, and where visitors can create associations within a museum collection, between collections, and between collections and the city. Those associations will form a resource for their later visits, but also for the visits of others.

Secondly, our current prototype system is being tailored and adapted to suit the Mack Room and to take account of the results of our studies and exploratory evaluations. Project members are now using the initial prototype along with other departmental colleagues, Lighthouse staff and a few members of the public. We consider it particularly important to have the sociologist and museum studies experts in our project use the system themselves. However, we are not yet undertaking formal or sustained evaluation or observation. Instead, we make minor changes to the system from day to day in the light of these informal short-term trials. This has led to work on reducing the delay between a 'wearable' visitor moving into a thematic zone and the automatic Web page display, refining the position tracking, and changing the design of the pages shown on the handheld so as to better combine map and content.

In the course of these early and small design iterations, we improve our understanding of the Room, its visitors and our technology. While we have studied the Mack Room and its visitors before, and our system before, the combination is new. We have many ideas that may be implemented in the near future, such as more explicit support for the visitors acting as guides, graphical representation of visitors' paths through the room, and recommendation of people, places and artifacts. Still, we wish to leave space for new ideas and priorities to grow from our ongoing evaluation and observation. In pausing before a major redesign, we wish to secure a balance between the studies of visitors, studies of visitors using our system, technological interests in city–wide use and adaptive information, and theoretical work on notions of social context, interpretation and media.

Conclusion

Social context is an essential factor of the museum experience. Social interaction with companion(s) and other members of the public directly influences a museum experience. In the light of new technologies that support remote access to museum settings, we can see that the social context of the

46

museum visit can extend beyond a physical room or building, and include computer-mediated interactions. Similarly, museums can enrich the digital site by better integrating it with the physical site. The *City* project aims at both understanding and supporting social interaction in cultural institutions; interaction that may be synchronous or asynchronous, and that may involve multiple media such as mobile computers, hypermedia and virtual environments. Our system supports interaction between physical and digital visitors, as well as between visitors and their environments, in a dynamic and contextual way. Our aim is a canvas of interaction that spans different places, people and times, with a focus on the creation and delivery of cultural information that is contextually appropriate, useful and engaging.

Acknowledgements

Our thanks go to Ian MacColl and Barry Brown (Glasgow), Cliff Randell (Bristol), Anthony Steed (UCL), Chris Greenhalgh, Ian Taylor & Tom Rodden (Nottingham), Dave Millard (Southampton), Dorothy McKay (The House for an Art Lover), Lynn Bennett and Stuart MacDonald (The Lighthouse). We gratefully acknowledge the EPSRC's funding, and the donation of equipment within Hewlett Packard's Art & Science programme.

References

Aoki, P. M. & A. Woodruff, (2000). Improving Electronic Guidebook Interfaces Using a Task-Oriented Design Approach. In *Designing Interactive Systems 2000*. ACM. 319-325.

Barbieri, T. & P. Paolini, (2000). Cooperative Visits for Museum WWW Sites a Year Later: Evaluating the Effect. In D. Bearman & J. Trant (Eds.) *Museums and the Web 2000*. Archives & Museum Informatics.

Chalmers, M. (2001). *Tales in the City: Adaptive Information in the Physical City*. Consulted February 11, 2002. (http://www.dcs.gla.ac.uk/equator/2001/01/vee.html).

Communications Design Team (1976). Spatial considerations. In Communications Design Team (Eds.) *Communicating with the Museum Visitors*. Royal Ontario Museum.

47

Diamond, J. (1986). The behaviour of Family Groups in Science Museums. *Curator*, 29(2). 139-154.

Falk, H. J. & L. D. Dierking, (1992). *The museum experience*. Washinghton: Whalesback Books.

Galani, A. (2001). *Vee in the Mackintosh Room*. Consulted February 11, 2002. (http://www.dcs.gla.ac.uk/equator/2001/05/scenario.html).

Lord, M. (1999). Editorial. *Museum International*, October 1999, Vol. 51 No 4. 3-3.

MacDonald, S. (1993). The enigma of the visitor sphinx. In S. Bicknell & G. Farmelo (Eds.) *Museum Visitor Studies in the 90s*. London: Science Museum, 77-85.

McManus, P.M. (1987). It's the Company you Keep... The Social Determination of Learning-related Behaviour in a Science Museum. *International Journal of Museum Management and Curatorship*, 6. 263-270.

Oppermann, R. & M. Specht, (1999). A Nomadic Information System for Adaptive Exhibition Guidance. In D. Bearman & J. Trant (Eds.) *Cultural Heritage Informatics: selected papers from ichim99*. Archives and Museum Informatics. 103-109.

Psarra, S. & T. Grajewski & M. O'Neil, (2002). Rethinking museum design – the effect of layout on visitors patterns of movement and use. *Museum Practice*, March 2002.

Randell, C. & H. Muller, (2001). Low Cost Indoor Positioning System. In G.D. Abowd, B. Brumitt & S. Shafer (Eds.) *Ubicomp 2001: Ubiquitous Computing*. Springer. 42-48.

Stenton, P. (2001). A user's view of the Equator Mackintosh Room Demo. Consulted February 11, 2002. (http://www.dcs.gla.ac.uk/equator/2001/06/stenton/stenton.html).

Vom Lehn, D., C. Heath & J. Hindmarsh, (2001). Exhibiting Interaction: Conduct of Collaboration in Museums and Galleries. *Symbolic Interaction*, Vol. 24 No 2. 189-216.

The Museum Wearable: real-time sensor-driven understanding of visitors' interests for personalized visually-augmented museum experiences

Flavia Sparacino, MIT Media Lab, USA

http://www.media.mit.edu/~flavia/projects.html

Abstract

This paper describes the museum wearable: a wearable computer which orchestrates an audiovisual narration as a function of the visitors' interests gathered from their physical path in the museum and length of stops. The wearable consists of a lightweight and small computer that people carry inside a shoulder pack. It offers an audiovisual augmentation of the surrounding environment using a small, lightweight eye-piece display (often called private-eye) attached to conventional headphones. Using custom built infrared location sensors distributed in the museum space, and statistical mathematical modeling, the museum wearable builds a progressively refined user model and uses it to deliver a personalized audiovisual narration to the visitor. This device will enrich and personalize the museum visit as a visual and auditory storyteller that is able to adapt its story to the audience's interests and guide the public through the path of the exhibit.

Keywords: wearable computer, handheld, personalized experience, Bayesian networks, location aware mobile computing

1. Introduction

In the last decade museums have been drawn into the orbit of the leisure industry and compete with other popular entertainment venues, such as cinemas or the theater, to attract families, tourists, children, students, specialists, or passersby in search of alternative and instructive entertaining experiences. Some people may go to the museum for mere curiosity, whereas others may be driven by the desire for a cultural experience. The museum visit can be an occasion for a social outing, or become an opportunity to meet new friends. While it is not possible to design an exhibit for all these categories of visitors, it is desirable to offer exhibit designers ways to attract as many as possible amongst the variety of individual visitors or visitor categories. Technology can help offer the museum public opportunities to personalize their visits according to their desires and expectations.

Traditional storytelling aids for museums have been signs and text labels, spread across the exhibit space; exhibit catalogues, typically sold at the museum store; guided tours, offered to groups or individuals; audio tours; and more recently video or multimedia kiosks with background information on the displayed objects. Each of these storytelling aids has advantages and disadvantages. Catalogues are usually attractive and well done, yet they are often too cumbersome to carry around during the visit as a means to offer guidance and explanations. Guided tours take away from visitors the choice of what they wish to see and for how long. They can be highly disruptive for the surrounding visitors, and their effectiveness strictly depends on the knowledge, competence, and communicative skills of the guide. Audio tours are a first step to help augment the visitor's knowledge. Yet when they are button activated, as opposed to having a location identification system, they can be distracting for the visitor. The information conveyed is also limited by the only-audio medium: it is not possible to compare the artwork described with previous relevant production of the author, nor to show other relevant images. Interactive kiosks are more frequently found today in museum galleries. Yet they are physically distant from the work they describe, thus do not support the opportunity for the visitors to see, compare, and verify the information received against the actual object. The author's experience suggests that when extensive Web sites are made available through

BEST COPY AVAILABLE

Museums and the Web 2002: Proceedings
Archives & Museum Informatics, 2002, p. 41

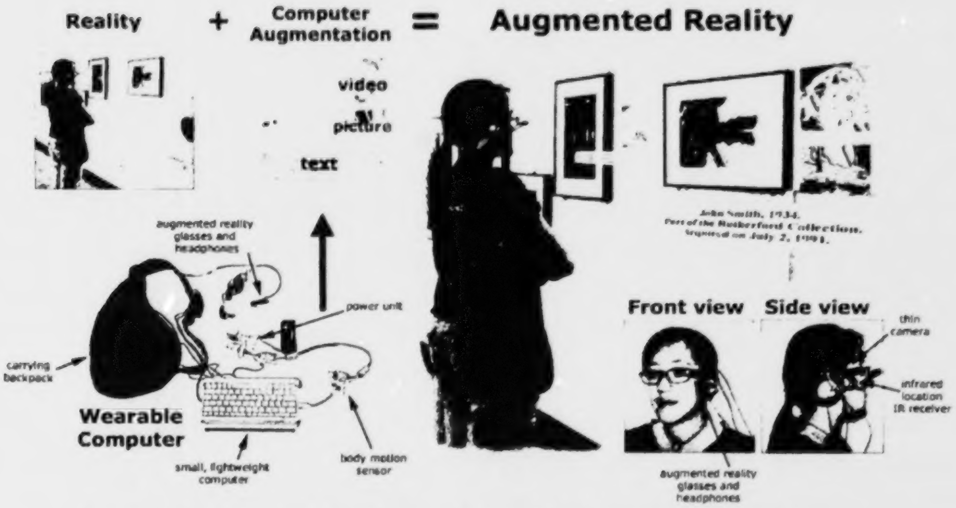
Reality
+
Computer Augmentation
=
Augmented Reality
video
picture
text
augmented reality glasses and headphones
power unit
carrying backpack
Wearable Computer
small, lightweight computer
body motion sensor
John Smith, 1934.
Part of the Rutherford Collection.
Acquired on July 2, 1991.
Front view
Side view
thin camera
infrared location IR receiver
augmented reality glasses and headphones

Figure 1. The Museum Wearable: explanation of concept and application

interactive kiosks placed along the museum galleries, these may absorb lengthy amounts of the visitors' museum time, thereby detracting from, rather than enriching, the objects on display. Finally, panels and labels with text placed along the visitors' path can interrupt the pace of the experience as they require a shift of attention from observing and contemplating to reading and understanding (Klein, 1986).

Indeed, when we walk through a museum there are so many different stories we could be told. Some of these are biographical about the author of an artwork; some are historical and allow us to comprehend the style or origin of the work; and some are specific about the artwork itself, in relation to other artistic movements. Museums usually have large Web sites with multiple links to text, photographs, and movie clips to describe their exhibits. Yet it would take hours for a visitor to explore all the information in a kiosk, to view the VHS cassette tape associated with the exhibit, and read the accompanying catalogue. Most people do not have the time to devote nor motivation to assimilate this type of information, and therefore the visit to a museum is often remembered as a collage of first impressions produced by the prominent feature of the exhibits, and the learning opportunity is missed. How can we tailor content to the visitors in a museum, during their visits, to enrich both the learning and entertainment experience ? We want a system which can be personalized to be able to dynamically create and update paths through a large database of content and deliver to the users in real time during the visit all the information they desire. If the visitors spend a lot of time looking at a Monet, the system needs to infer that the users likes Monet, and update the paths through the content to take that into account. This research proposes the museum wearable as an effective way to turn this scenario into reality.

This document illustrates the hardware, authoring techniques, and software created for the construction of the museum wearable. It speculates on its ability to facilitate a new style of exhibit design, and postponesits, as future work, the assessment at the exhibit's site, of the museum wearable's contribution to the public's experience.

2. The Museum Wearable

Wearable computers have risen to the attention of technological and scientific investigation (Starner, 1997) and offer an opportunity to "augment" for visitors their perception/memory/experience of the exhibit in a personalized way. The museum wearable is a wearable computer which orchestrates an audiovisual narration as a function of the visitors' interests gathered from their physical path in the museum and length of stops. It offers a new type of entertaining and informative museum experience,

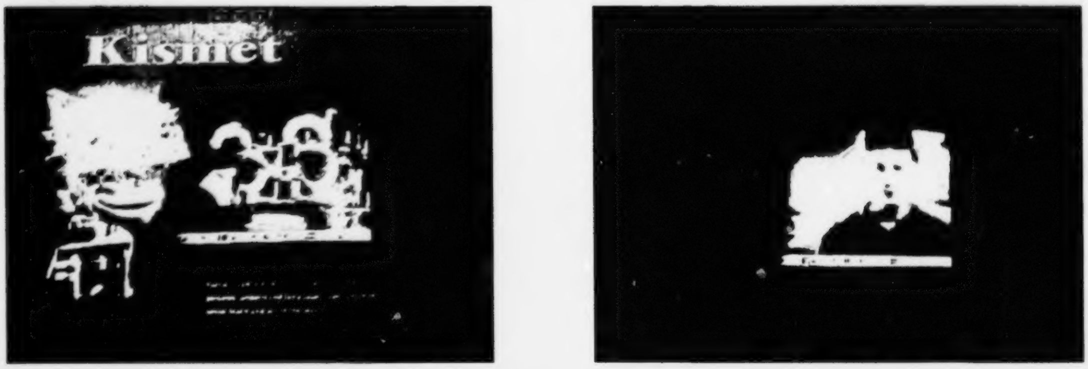

Figures 2, 3. Camera "wearing" the head mounted display: *shows how the user's brain assembles the real world's image seen by the unencumbered eye with the display's image seen by the other eye, into a fused augmented reality image.*

more similar to mobile immersive cinema than to the traditional museum experience [figure 1].

The museum wearable is made by a lightweight CPU hosted inside a small shoulder pack and a small, lightweight private-eye display. The display is a commercial monocular, VGA-resolution, color, clip-on screen attached to a pair of sturdy headphones. When wearing the display, after a few seconds of adaptation, the user's brain assembles the real world's image, seen by the unencumbered eye, with the display's image seen by the other eye, into a fused augmented reality image [figures 2, 3].

The wearable relies on a custom-designed long-range infrared location-identification sensor to gather information on where and how long the visitor stops in the museum galleries. A custom system had to be built for this project to overcome limitations of commercially available infrared location identification systems such as short range and narrow cone emission. The location system is made by a network of small infrared devices, which transmit a location identification code to the receiver worn by the user and attached to the display glasses. The transmitters have the size of a 9V battery, and are placed inside the museum, next to the regular museum lights. They are built around a microcontroller, and their signal can be reliably detected at least as far as 30 feet away within a large cone range of approximately ten to thirty degrees, according to the area that needs to be covered, and up to 100 feet along a straight line. The emitter location identification tags have been embedded inside standard light fixtures to allow the exhibit designer to easily place them in the museum, next to the regular museum lights, and using the same power rack as the regular museum spotlights [figures 4, 5, 6, 7].

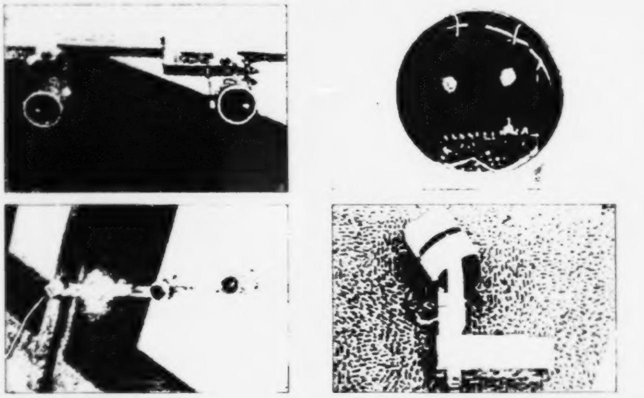

Figures 4, 5, 6, 7. Location sensor: emitter tags embedded inside light fixtures

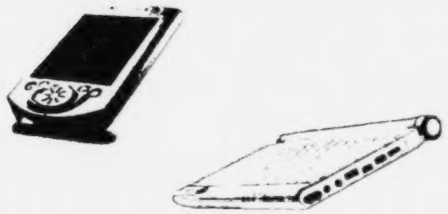

Figure 8. PAQ pocket PC 3670 and Figure 9. Sony picturebook with removed LCD

In view of having a museum wearable which can later be expanded to include other sensors, and process information not just from the infrared location sensor, but for example also from a small camera processing images in real time, a commercially available processing unit has been selected for this project. The processor of choice is a small-sized laptop computer: the SONY picturebook PCG-C1VPK, selected for its combined size, weight, computing power, multimedia capabilities, and long-lasting batteries. Given that the images generated by the laptop are viewed uniquely through the head mounted display, the LCD screen has been removed from the picturebook, to reduce weight and size. The picturebook features a Crusoe™ processor TM5600 clocked at 667 MHz, and without the LCD weighs only approximately one lb, and has a size of 0.5" X 9.8" X 6.0" (H x W x D). The picturebook has a 15GB capacity hard drive, which allows the programmer to store on the local hard drive several hours (8-10) of MPEG-compressed VGA resolution video (640x480) (approximately one hour of MPEG-compressed 640x480 video per one GB of available space on the internal hard drive). It also has 128 MB SDRAM, which allows the computer to play smoothly audio and video clips, as well as process images in real time when the computer is connected to a camera. The external ports include one USB port which is connected to the infrared receiver with a USB to serial converter, and a VGA and headphone output, which are connected to the video/audio inputs of the head-mounted display of the museum wearable. It also supports one type II card, which can be used to host a PCMCIA card for wireless communication over the Internet or a PCMCIA video acquisition card. All these features, combined with a battery life of 2.5-5.5 hours with the standard lightweight battery, large enough for a single museum visit, make the picturebook an ideal choice for the selected application [diagram 1].

An alternative to the picturebook is the smaller handheld IPAQ pocket PC 3670 [figures 8, 9]. The iPAQ 3670 features 64 MB of SDRAM and a 206-MHz Intel StrongARM SA-1110 32-bit RISC Processor. It has USB or serial connectivity that would interface with the infrared receiver of the museum wearable and it is only 5.11" x 3.28" x 0.62" (HxWxD) in size. The iPAQ 3670 is a desirable solution for the deployment of several museum wearables – which need to work only with the location sensor – in a museum. The iPAQ solution is cost effective because its price, with the necessary accessories, is about half the price of the SONY picturebook, and its size is also smaller.

The size and weight of both the wearable's processor and augmented reality display are critical for a museum application. The display cannot have a heavy and power hungry powering unit which requires frequent battery changes. Glasses also need to easily fit various people's head sizes, with annexed hair style, which is not an easy task. The wearable would be handed out to between ten and one hundred people a day, and therefore needs to be of robust assembly and easy to wear. Two different design solutions were implemented after a thorough iterative design process which included considerations about video resolution, power consumption, purchase availability and foremost, ergonomics, with

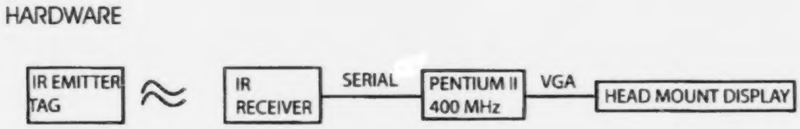

Diagram 1: The museum wearable hardware architecture

Figures 10,11. The common fashion and high fashion private eye display

the inclusion of adaptability with sensors, weight, and size. The common fashion display features an augmented reality display, which joins together a lightweight VGA resolution color display from the MicroOptical corporation, and a commercial high quality sturdy set of headphones [figure 10]. The high fashion display is a provocative stylish mount, mainly intended for visitors with a strong sense of aesthetics, and suitable for use in wearable fashion shows, to promote a non-nerdy and highly fashionable wear of augmented reality displays. With this design the MicroOptical augmented reality display is rigidly mounted to a pair of Oakley "over the top" glasses, as illustrated in figure 11.

Summary of elements of the museum wearable hardware:

- containing shoulder pack
- computer (CPU): SONY picturebook from which the display has been removed to reduce weight
- Head Mounted Display (HMD): VGA resolution MicroOptical clip-on mounted on sturdy headphones with a custom mount
- HMD's powering unit: hosted inside the containing shoulder pack
- Infrared receiver: the sensor is located on top of the headphones and the receiver circuit is located inside the containing shoulder pack.

The museum wearable plays an interactive audiovisual documentary about the displayed artwork on the private-eye display. Each mini-documentary is made of small segments which vary in size from twenty seconds to one and half minutes. A video server, written in C++ and DirectX-8, plays these assembled clips and receives TCP/IP messages from another program containing the information measured by the location ID sensors. This server-client architecture allows the programmer to easily add other client programs to the application, to communicate to the server information from other possible sources such as sensors or cameras placed along the museum aisles to measure how crowded the galleries are or how often a certain object has been visited. The client program reads IR data from the serial port, and the server program does inference, content selection, and content display using DirectX for full screen play back of the MPEG compressed clips [diagram 2].

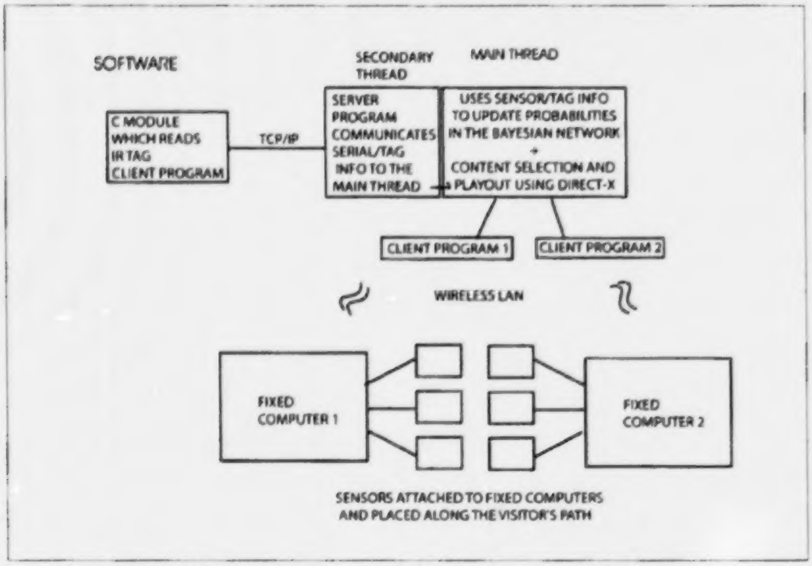
SOFTWARE
SECONDARY THREAD
MAIN THREAD
C MODULE WHICH READS IR TAG CLIENT PROGRAM
TCP/IP
SERVER PROGRAM COMMUNICATES SERIAL/TAG INFO TO THE MAIN THREAD
USES SENSOR/TAG INFO TO UPDATE PROBABILITIES IN THE BAYESIAN NETWORK
+
CONTENT SELECTION AND PLAYOUT USING DIRECT-X
CLIENT PROGRAM 1
CLIENT PROGRAM 2
WIRELESS LAN
FIXED COMPUTER 1
FIXED COMPUTER 2
SENSORS ATTACHED TO FIXED COMPUTERS AND PLACED ALONG THE VISITOR'S PATH

Diagram 2: The museum wearable software architecture

3. Visitors' interests and visitor types

To identify visitor preferences, museums seek to identify target groups: individuals who share common traits such as culture, ethnic or social affiliation, educational level, and leisure preferences. Curators and designers need to assess the basic knowledge and expectations of these groups to be able to reach, communicate with, and stimulate curiosity in all of their visitors. Eleanor Hooper-Greenhill identifies target groups which include families, school parties, other organized educational groups, leisure learners, tourists, the elderly, and people with visual, auditory, mobility or learning disabilities (Hooper-Greenhill, 1999, p. 86). She then suggests a partition of museum resources, to target, attract and entertain these different groups. During a personal interview, Beryl Rosenthal, director of exhibitions at the MIT Museum, described a more sophisticated visitor type classification. She identified stroller moms, accompanied by children three years old or younger; window shoppers: families who cruise through the museum in search of an alternative leisure experience; button pushers, typically adolescents; school groups; the date crowd; and the PhDs, who want to know (and criticize) everything in the museum. Young visitors, children 5-14 also represent a separate group of visitors with different learning needs and curiosities than the other groups. While this colorful classification well depicts the variety of public that museums need to attract, entice and educate, it is too sophisticated to model mathematically, at least initially.

More usefully for this research, Dean generalizes museum visitors in three broad and much simpler categories (Dean, 1994, pp. 25-26). The first category includes what he calls the "casual visitors": people who move through a gallery quickly and who do not become heavily involved in what they see. Casual visitors use some of their leisure time in museums but do not have a strong stimulus or motivation to deepen their knowledge about the objects on display. The second group, the "cursory visitors," show instead a more genuine interest in the museum experience and their collections. According to Dean, these visitors respond strongly to specific objects that stimulate their curiosity and wander through the gallery in search of further such stimulus for a closer exploration of the targeted objects. They do not read every label nor absorb all available information, but will occasionally read and spend time in selected areas or with selected objects of interest they encounter in the galleries. The third group is a minority of visitors who thoroughly examine exhibitions with much more detail and attention. They are learners who will spend an abundance of time in galleries, read the text and labels, and closely examine the objects. Dean attributes differences between "people who rush", "people who stroll", and "people who study" to different prior experiences and educational level. Yet he states that it is important for museums to be equipped to communicate with and interest all visitors by scaling and designing an exhibit so that it offers entertainment to the "stroller" as well as an opportunity to deepen knowledge for the "learners".

Serrell (1996) also divides visitors into three types: the transient, the sampler and the methodical viewers. She notes that currently museum evaluators are using terms like "streakers, studiers, browsers, grazers and discoverers" to characterize museum visitors' styles of looking at exhibits. But she concludes that this type of categorization is not useful for summative evaluation, suggesting that it is a subjective method of classification, and that it is not fruitful to try to create exhibitions that serve these different styles of visiting. Instead she suggests that a more objective means of classification be found, such as average time spent in the exhibition space.

In accordance with the simplified museum visitor typology suggested by Dean and Serrell, the museum wearable identifies three main visitor types. To offer a more intuitive understanding of these types they have been renamed: the busy, selective, and greedy visitor type. The greedy type wants to know and see as much as possible, and does not have a time constraint; the busy type just wants to get an overview of the principal items in the exhibit, and see little of everything; and the selective type wants to see and know in depth only about a few preferred items. The identification of other visitor types or subtypes has been postponed to future improvements and developments of this research.

The visitor type estimation is obtained probabilistically with a Bayesian network using as input the information provided by the location identification sensors on where and how long the visitor stops, as if the system were an invisible storyteller following the visitor in the galleries and trying to guess preferences based on observation of external behavior.

4. Sto(ry)chastics: Sensor-driven Understanding of Visitors' Interests with Bayesian Networks

In order to deliver a dynamically changing and personalized content presentation with the museum wearable, a new content authoring technique had to be designed and implemented. This called for an alternative method than the traditional complex centralized interactive entertainment systems which simply read sensor inputs and map them to actions on the screen. Interactive storytelling with such one-to-one mappings leads to complicated control programs which have to do an accounting of all the available content, where it is located on the display, and what needs to happen when/if/unless. These systems rigidly define the interaction modality with the public, as a consequence of their internal architecture, and lead to presentations which have little depth of content, are hard to modify, ad hoc, and prone to error. The main problem with such content authoring approaches is that they acquire high complexity when drawing content from a large database, and that once built, they are hard to modify or to expand upon. In addition, when they are sensor-driven they become depended on the noisy sensor measurements, which can lead to errors and misinterpretation of the user input (Sparacino, 1999). Rather than directly mapping inputs to outputs, the system should be able to "understand the user" and to produce an output based on the interpretation of the user's intention in context.

To overcome the limitations of one-to-one mapping systems, the mathematical modeling approach developed to author content for the museum wearable is a real-time sensor-driven stochastic modeling of story and user-story interaction. It has therefore been called sto(ry)chastics. Sto(ry)chastics uses graphical probabilistic modeling of story fragments and participant input, gathered from sensors, to tell a story to the user, as a function of estimated intentions and desires during interaction. With this approach, the coarse and noisy sensor inputs – path and length of stops given by the location sensors – are coupled to digital media outputs via a user model, estimated probabilistically by a Bayesian network.

A Bayesian network is a graphical model which encodes probabilistic relationships amongst variables

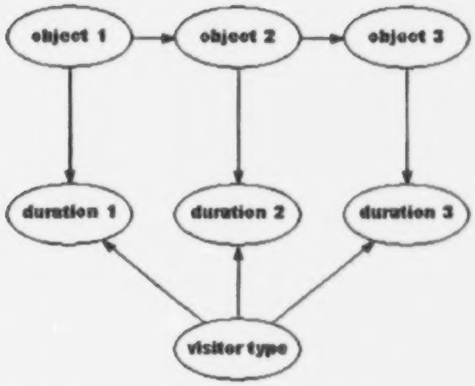
object 1
object 2
object 3
duration 1
duration 2
duration 3
visitor type

Figure 12. Topology of the Bayesian network for visitor time estimation, limited to three objects/time slices.

of interest (Pearl, 1988; Jordan, 1999). Such graphs not only provide an attractive means for modeling and communicating complex structures, but also form the basis for efficient algorithms, both for propagating evidence and for learning about parameters. Bayesian networks encode qualitative influences between variables in addition to the numerical parameters of the probability distribution. As such they provide an ideal form for combining prior knowledge and data. By using graphs, not only does it become easy to encode the probability independence relations amongst variables of the network, but it is also easy to communicate and explain what the network attempts to model (Smyth, 1997). Graphs are easy for humans to read, and they help focus attention, for example in a group working together to build a model. This allows the digital architect, or the engineer, to communicate on the same ground (the graph of the model) as the curator and therefore to be able to encapsulate the curator's domain knowledge in the network, together with the sensor data.

Sto(ry)chastics uses a Bayesian network to estimate the user's preferences taking the location identification sensor data as the input or observations of the network. The user model is progressively refined as the visitor progresses along the museum galleries: the model is more accurate as it gathers more observations about the user. Figure 12 shows the Bayesian network for visitor estimation, limited to three museum objects (so that the figure can fit in the document), selected from a variety of possible models designed and evaluated for this research. Figures 13 and 14 show state values for the network after the visitor has made a long stop at the first object, followed by another long stop at the second object.

The sto(ry)chastics approach has several advantages. It is:

1. **Flexible:** it is possible to easily test many different scenarios by simply changing the parameters of the system.

2. **Reconfigurable:** it is also quite easy to add or remove nodes and/or edges from the network without having to "start all over again" and specify again all the parameters of the network from scratch. This is a considerable and important ad-

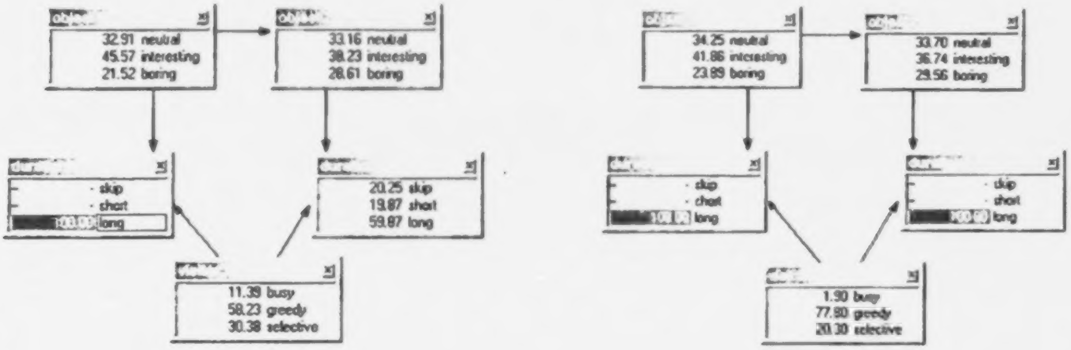
32.91 neutral
45.57 interesting
21.52 boring
33.16 neutral
38.23 interesting
28.61 boring
skip
short
100.00 long
20.25 skip
19.87 short
59.87 long
11.39 busy
58.23 greedy
30.38 selective
34.25 neutral
41.86 interesting
23.89 boring
33.70 neutral
36.74 interesting
29.56 boring
skip
short
long
skip
short
long
1.90 busy
77.80 greedy
20.30 selective

Figures 13, 14. Values of the network nodes after the visitor has made a long stop at the first object followed by another long stop at the second object.

BEST COPY AVAILABLE

vantage with respect to hard coded or heuristic approaches to user modeling and content selection. Only the parameters of the new nodes and the nodes corresponding to the new links need to be given. The system is extensible story-wise and sensor-wise. These two properties: flexibility and ease of model reconfiguration, allow the system engineer, the content designer, and the exhibit curator to work together and easily and cheaply try out various solutions and possibilities until they converge on a model which satisfies all the requirements and constraints for their project. A network can also rapidly be reconfigured for another exhibit.

3. **Robust:** Probabilistic modeling allows the system to achieve robustness with respect to the coarse and noisy sensor data.

4. **Adaptive:** sto(ry)chastics is adaptive in two ways: it adapts both to individual users and to the ensemble of visitors of a particular exhibit. For individuals, even if the visitor exhibits an initial "greedy" behavior, it can later adapt to the visitor's change of behavior. It is important to notice that, reasonably and appropriately, the system "changes its mind" about the user type with some inertia: i.e. it will initially lower the probability for a greedy type until other types gain probability. Sto(ry)chastics can also adapt to the collective body of its users. If a count of busy/greedy/selective visitors is kept for the exhibit, these numbers can later become priors of the corresponding nodes of the network, thereby causing the entire exhibit to adapt to the collective body of its users through time. This feature can be seen as "collective intelligence" for an exhibit which can adapt not just to the individual visitors but also to the set of its visitors.

5. **Context-sensitive:** for any system to be robust and to provide relevant information to its user, it is important to model the context of interaction together with the other system parameters. For example sto(ry)chastics could provide an explanation of a visitor's change of behavior at the museum. If suddenly a greedy type starts making short stops, the system, before concluding that that visitor is actually a selective or busy type, could test if the current time is near closing time for the museum galleries, or if, by use of other room sensors, there is a great crowd in the galleries where the visitor is making short stops. Coming up with the right conclusions, given this type of external information, means that the system is context-sensitive.

Sto(ry)chastics has therefore implications both for the human author (designer/curator) who is given a flexible modeling tool to organize, select, and deliver the story material, as well as the audience, who receives personalized content only when and where it is appropriate.

5. Experimentation: the Museum Wearable at the MIT Museum's Robots and Beyond Exhibit

The ongoing robotics exhibit at the MIT Museum provided an excellent platform for experimentation and testing with the museum wearable [figures 15, 16]. This exhibit, called Robots and Beyond, and curated by Janis Sacco and Beryl Rosenthal, features landmarks of MIT's contribution to the field of robotics and Artificial Intelligence. The exhibit is organized in five sections: Introduction, Sensing, Moving, Socializing, and Reasoning and Learning, each including robots, a video station, and posters with text and photographs which narrate the history of robotics

Figures 15 and 16. Visitor with the Museum Wearable at MIT Museum's Robots and Beyond Exhibit

BEST COPY AVAILABLE

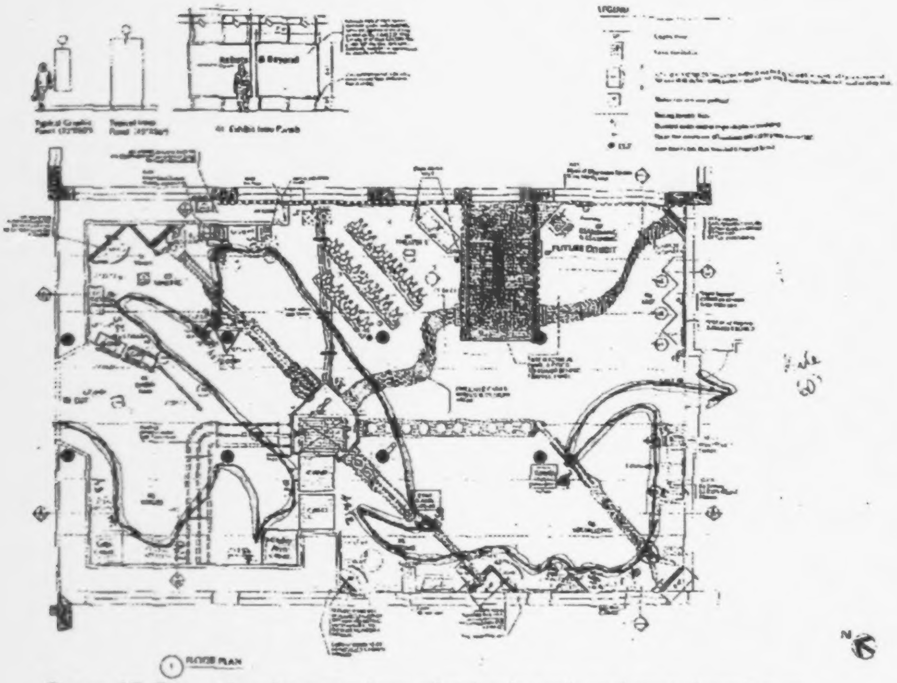

Figure 17. Example of annotations of visitor's path and stop duration at MIT Museum's Robots and Beyond exhibit

at MIT. There is also a large general purpose video station with large benches for people to have a seated stop and watch a PBS documentary featuring robotics research from various academic institutions in the country.

In order to set the initial values of the parameters of the Bayesian network, experimental data was gathered on the visitors' behavior at the Robots and Beyond exhibit. According to the VSA (Visitor Studies Association, http://museum.cl.msu.edu/vsa), timing and tracking observations of visitors are often used to provide an objective and quantitative account of how visitors behave and react to exhibition components. This type of observational data suggests the range of visitor behaviors occurring in an exhibition, and indicates which components attract, as well as hold, visitors' attention (in the case of a complete exhibit evaluation this data is usually accompanied by interviews with visitors, before and after the visit).

During the course of several days a team of collaborators tracked and made annotations about the visitors at the MIT Museum. Each member of the tracking team had a map and a stop watch. Their task was to draw on the map the path of individual visitors, and annotate the locations at which visitors stopped, the object they were observing, and for how long they stopped. In addition to the tracking information, the team of evaluators was asked to assign a label to the overall behavior of the visitor, according to the three visitor categories earlier described: "busy", "greedy", or "selective" [figure 17].

A subset of twelve representative objects of the Robots and Beyond exhibit were selected to evaluate this research, to shorten editing time. The geography of the exhibit needed to be reflected into the topology of the network, as shown in figure 18. Additional objects/nodes of the modeling network can be added later for an actual large scale installation and further revisions of this research.

The visitor tracking data is used to learn the parameters of the Bayesian network. The model can later be refined; that is, the parameters can be fine-tuned, as more visitors experience the exhibit with the museum wearable. The network has been tested and validated on this observed visitor tracking data

BEST COPY AVAILABL

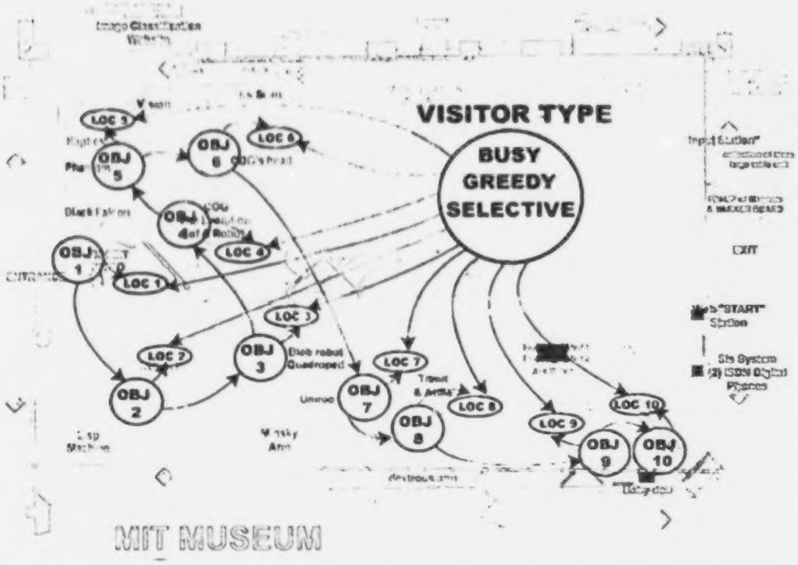

Figure 18. The geography of the exhibit needs to be reflected into the topology of the network

by parameter learning using the Expectation Maximization (EM) algorithm, and by performance analysis of the model with the learned parameters, with a recognition rate of 0.987 (Sparacino, 2001).

The system works in two steps. The first is user type estimation as described above. The next step is to assemble a mini-story for the visitor, relative to the nearest object [figures 19, 20, 21]. Most of the audio-visual material available for use by the museum wearable tends to fall under a set of characterizing topics, which typically define art and science documentaries. This same approach to documentary as a composition of segments belonging to different themes has been developed by Houbart in her work which edits a documentary based on the viewer's theme preferences, as an offline process (Houbart, 1994). The difference between Houbard's work and what the museum wearable does is that the museum wearable performs editing in real time, using sensor input and Bayesian network modeling to figure out the user's preferences (type). After an overview of the audio-visual material available at MIT's Robots and Beyond exhibit, the following content labels, or bins, were identified to classify the component video clips:

- **Description of the artwork:** what it is, when it was created (answers: when, where, what)
- **Biography of author:** anecdotes, important people in artist's life (answers: who)
- **History of the artwork:** previous relevant work of the artist
- **Context:** historical, what is happening in the world at the time of creation
- **Process:** particular techniques used or invented to create the artwork (answers: how)
- **Principle:** philosophy or school of thought the author believes in when creating the artwork (answers: why)
- **Form and Function:** relevant style, form and function which contribute to explain the artwork.
- **Relationships:** how is the artwork related to other artwork on display
- **Impact:** the critics' and the public's reaction to the artwork

This project required a great amount of editing to be done by hand (non automatically) in order to segment the two hours of video material available for the Robots and Beyond Exhibit at the MIT museum in the smallest possible complete segments. After this phase, all the component video clips were given a name, their length in seconds was recorded into the system, and they were also classified according to the list of bins described above. The classification was done probabilistically; that is, each clip has been assigned a probability (a value between

BEST COPY AVAILABLE

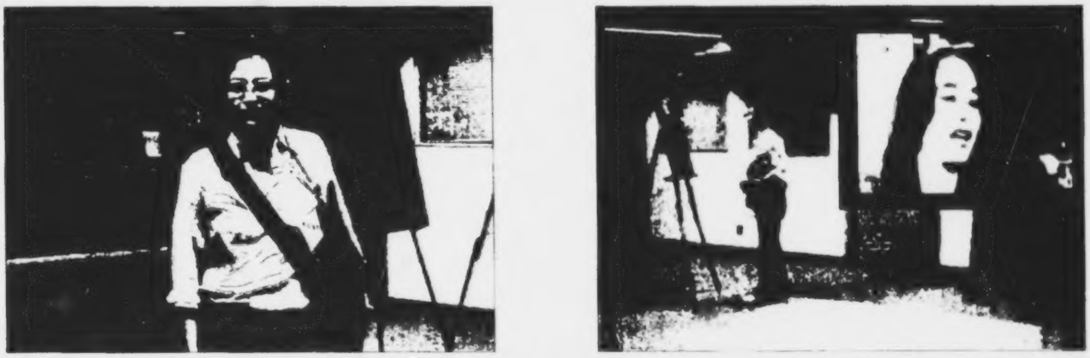

Figures 19, 20. Visitor wearing the museum wearable and receiving an audiovisual story about the displayed artwork (picture in picture).

zero and one) of belonging to a story category. The sum of such probabilities for each clip needs to be one. The result of the clip classification procedure, for a subset of available clips, is shown in table 1.

To perform content selection, "conditioned" on the knowledge of the visitor type, the system needs to be given a list of available clips and the criteria for selection. There are two competing criteria: one is given by the total length of the edited story for each object, and the other is given by the ordering of the selected clips. The order of story segments guarantees that the curator's message is correctly passed on to the visitor, and that the story is a "good story", in that it respects basic cause-effect relationships and makes sense to humans. Therefore the Bayesian network described in the previous paragraph needs to be extended with additional nodes for content selection [figures 22, 23]. The additional "good story" node, encodes, as prior probabilities, the curator's preferences about how the story for each object should be told.

A study of how content is distributed geographically along the exhibit, both in two and three dimensions, was also performed. The purpose of this

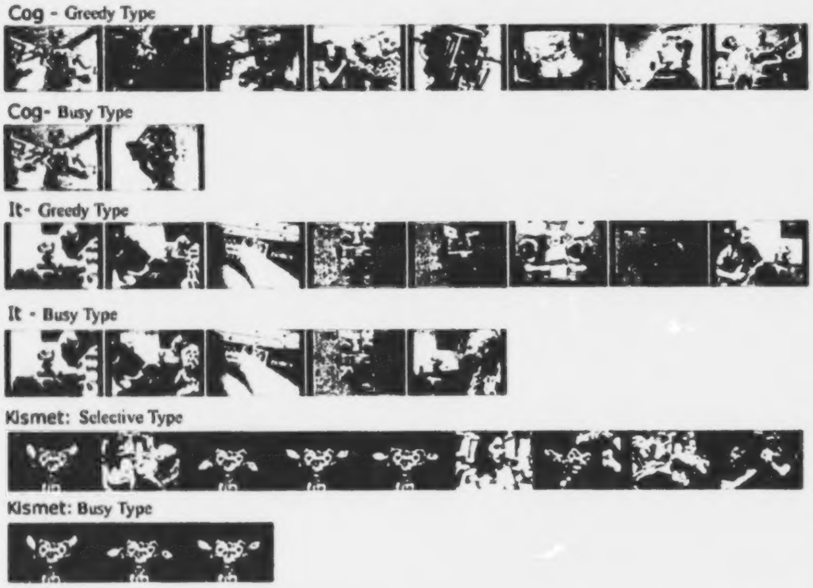

Figure 21. Storyboards from various video clips shown on the museum wearable's display at MIT Museum's Robots and Beyond Exhibit

CATEGORIES / TITLES LENGTH IN SECS		[illegible] 021	[illegible] 090	Cogdrum 063	Cogfuture 043	Coghistory 031
Description	DSC	0.7	0.2	0.3	0	0
History	HST	0.1	0	0	0	0.5
Context	CTX	0	0	0	0	0
Biography	BIO	0.1	0.2	0	0	0.1
Process	PRC	0	0	0.6	0	0
Principle	PNC	0	0.4	0.1	1	0.2
Form & Function	FAF	0	0.2	0	0	0.2
Relationships	REL	0.1	0	0	0	0
Impact	IMP	0	0	0	0	0
Total P		1	1	1	1	1

CATEGORIES / TITLES LENGTH IN SECS		Cogintro 041	Dexdesign 114	[illegible] 034	Dexintro 072	[illegible] 096
Description	DSC	0.8	0.1	0.3	0.5	0.4
History	HST	0	0	0	0	0
Context	CTX	0	0	0	0.2	0
Biography	BIO	0	0.2	0	0	0
Process	PRC	0.2	0.3	0	0	0.2
Principle	PNC	0	0	0.2	0.1	0.4
Form & Function	FAF	0	0.4	0.5	0	0
Relationships	REL	0	0	0	0	0
Impact	IMP	0	0	0	0.2	0
Total P		1	1	1	1	1

CATEGORIES / TITLES LENGTH IN SECS		[illegible] 026	[illegible] 057	[illegible] 066	[illegible] 160	[illegible] 066
Description	DSC	0.8	0	0.1	0.3	0.3
History	HST	0	0	0.2	0	0
Context	CTX	0	0	0	0	0
Biography	BIO	0.2	0.9	0	0	0
Process	PRC	0	0	0.1	0.5	0.3
Principle	PNC	0	0	0.2	0	0.3
Form & Function	FAF	0	0	0.4	0.2	0.1
Relationships	REL	0	0.1	0	0	0
Impact	IMP	0	0	0	0	0
Total P		1	1	1	1	1

CATEGORIES / TITLES LENGTH IN SECS		[illegible] 183	[illegible] 255	[illegible] 021	Legflamingo 063	[illegible] 116
Description	DSC	0.1	0.3	1	0.5	0.4
History	HST	0	0	0	0	0
Context	CTX	0.1	0	0	0	0
Biography	BIO	0	0	0	0	0.2
Process	PRC	0.3	0.6	0	0.5	0.3
Principle	PNC	0.5	0	0	0	0.1
Form & Function	FAF	0	0	0	0	0
Relationships	REL	0	0.1	0	0	0
Impact	IMP	0	0	0	0	0
Total P		1	1	1	1	1

Table 1. Selected segments cut from the video documentation available for the MIT Museum's Robots and Beyond Exhibit. All segments have been assigned a set of probabilities which express their relevance with respect to nine relevant story themes or categories.

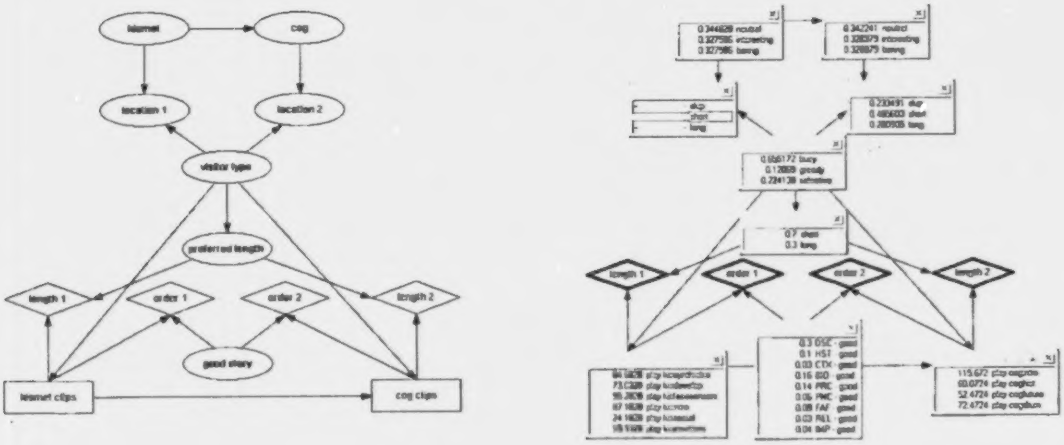
kismet
cog
location 1
location 2
visitor type
preferred length
length 1
order 1
order 2
length 2
good story
kismet clips
cog clips
skip
short
long
length 1
order 1
order 2
length 2

Figures 22, 23. Extension of Bayesian network to perform content selection.

BEST COPY AVAILABLE

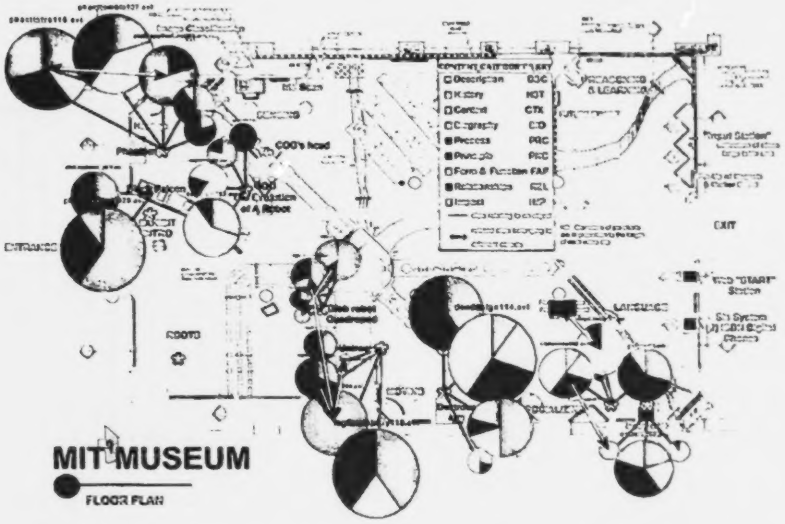

Figure 24. Two dimensional representation of content distribution for MIT Museum's Robots and Beyond Exhibit.

study was to visualize the different stories for different visitors edited by the museum wearable as paths through the hyperspace of content in the exhibit. The 2D study shows colored pie charts in the vicinity of the twelve tracked objects at the museum. Each pie chart represents the content available for the corresponding object. The size of the pie chart is proportional to the amount of content available for that object. The size of the colored slices of the chart represents the contribution of each story bin to the content available for the object [figure 24]. The 3D content map provides a visualization of how the content bins contribute to create a storyscape specific to this exhibit. It contains color coded vertical columns (a color for each content bin) whose height is proportional to the amount of content that each bin contributes to for the corresponding object [figures 25, 26].

For content personalization the system should be able to infer an interest profile of the visitors, in addition to their type as they wander along the exhibit gallery. An interest profile in the context of this research means a rating of preference for the story themes given in table 1. A GSR (Galvanic Skin Response) sensor can potentially give very useful information to the museum wearable. This sensor responds to skin conductivity and is often used in the medical and psychological field as an aid to monitor an individual's level of excitement or stress (Healey, 1999). If for example the GSR sensor measures a train of peaks when the wearable is playing a segment with biographical information about the portrayed artist, the system can infer, with a certain probability, that the visitor has a strong interest in this topic i.e. biography. It will then update the visitor interest profile with the gathered visitor preferences. The probabilistic framework offered by the Bayesian network approach is particularly relevant for this type of sensor. For example, the sensor could measure excitability for other reasons than that a compelling video segment being shown, such as meeting a friend, or recalling something that happened earlier during the day. These "false positive" data points would be modeled as "noise" intrinsic in the GSR measurements. The additional GSR sensor has been tested in a computer-based simulated environment. In this case, the decision node for content selection also needs to take into account the visitor's preferences, which compete with the curator's ordering preferences to assemble the best matching audiovisual story for each object.

6. Potential Impact of the Museum Wearable on Exhibit Design

Potential changes and improvements that the museum wearable can produce for the space design of the MIT Museum's Robots and Beyond exhibit became obvious in the course of this project.

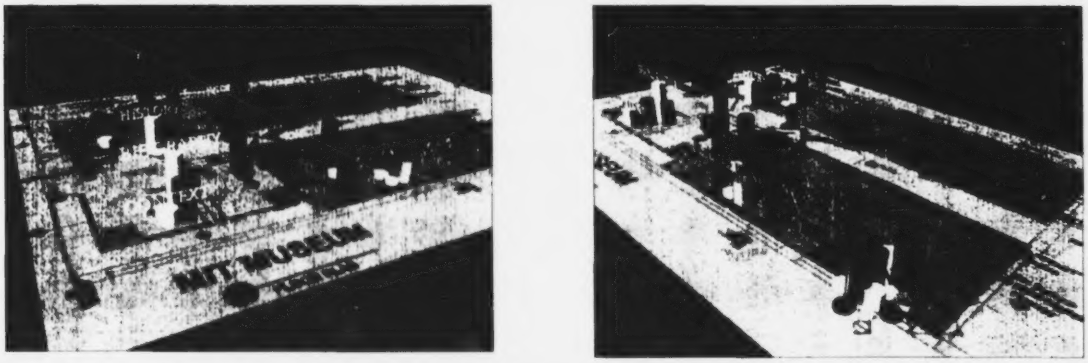

Figures 25, 26. Three dimensional representation of content distribution for MIT Museum's Robots and Beyond Exhibit.

At the current exhibit the posters and labels occupy half of the available exhibit wall space, and while they certainly provide useful information, they require long stops for reading, take useful space away from other interesting objects which could be displayed in their stead, and are not nearly as compelling and entertaining as a human narrator (a museum guide) or a video documentary about the displayed artwork. The tracking data and our observation of museum visitors also revealed that people do not spend sufficient time to read all of what is described in the posters to absorb the corresponding information. A great deal of the space occupied by the posters and text labels is therefore wasted, as most people don't take advantage of information provided in a textual form.

The video stations, located in each section of the exhibit, complete the narration about the artwork by showing the robots in motion and by featuring interviews with their creators. While the video stations provide compelling narrative segments, they are not always located next to the object described, and therefore the visitor needs to spend some time locating the described objects in the surrounding space in order to associate the object to the corresponding narrative segment. The video stations detract attention from the actual objects on display, and are so much the center of attention in the exhibit that the displayed objects seem to be more of a decoration around the video stations than the actual exhibit.

The potential improvements to the exhibit layout offered by the museum wearable are summarized as follows:

1. There would be no more need to have so many posters and text labels, as the corresponding information could be provided in a more appealing audio visual form, in a video documentary style by the museum wearable. The space now made available by eliminating the large posters could be used to display more robots, which are the true protagonists of the exhibit. Typically most exhibits have to discard many interesting objects as there is not enough physical space available in the museum galleries for all objects. Therefore making more space available is a clear advantage provided to the exhibit designer and the curator. Figures 27, 28 show how the posters at the entrance of the MIT Museum's Robots and Beyond exhibit can be replaced by more objects to be seen and appreciated by the public. These images are extracted from a three dimensional Alias Wavefront's Maya 3 animation realized to visualize the wearables' potential impact on the exhibit's space [figure 29].

2. Visitors would be better informed, as the information currently provided by the posters is mostly neglected by the public. The same information would instead become part of the overall narration provided by the wearable, and it would be better absorbed and appreciated by the public.

3. The video kiosks would no longer be necessary because the same material would be presented by the museum wearable. The robots would be again the center of attention for visitors, as the wearables allows both the real world and the augmented audiovisual information to be seen at

62

BEST COPY AVAILABLE

Figures 27, 28. Potential impact of the museum wearable on the current exhibit layout: the posters in the Roots section are replaced by new objects.

the same time as part of the wearer's real surround view. This would again make more space available for additional objects to be displayed.

The fact that the museum wearable presents audio visual material together with the corresponding object, rather than separately in space and time, and within the same field of view of the visitor, thanks to the private-eye display, is also of great importance. While no studies have been conducted yet on the quality and effectiveness of the learning experience offered by the museum wearable, there is reasonable hope that synchronous and local information provided while actually looking at the object described by the wearable can make a longer and more effective impression on the visitor. With this device curators may be able to present a larger variety of more connected material in an engaging manner within the limited physical space available for the exhibit.

7. Related Work

Oliver (http://www.media.mit.edu/~nuria/dypers/dypers.html; Schiele, 1999) developed a wearable computer with a visual input as a visual memory aid for a variety of tasks, including medical, training, or education. This system records small chunks of video of a curator describing a work of art, and associates them with triggering objects. When the objects are seen again at a later moment, the video is played back. The museum wearable differs from the previous application in many ways. DYPERS is a personal annotation device, and as opposed to the museum wearable, it does not attempt to perform either user modeling or a more sophisticated form of content selection and authoring. It does one-to-one associations between triggering objects and recording or play back of clips. Besides general training, is used specifically in the museum context to allow a visitor to record salient moments of the explanation by a human guide to later replay them in the context of an independent visit to a museum, without a guide. The museum wearable, in contrast, focuses on estimating the visitor's type and interest profile to deliver a flexible user-tailored narrative experience from audio/video clips that have been prerecorded. These clips or animations would usually be part of the museum's digital media collection. As opposed to DYPERS, it does not have the ability to record new content to be played out at a later time. Its purpose is to create for the visitor a path-driven personalized and immersive cinematic experience which takes into account the overall trajectory of the visitor in the museum, the amount of time that visitor stays to look at and explore the objects on display, to select a personalized story for the visitor, out of several possible digital stories that can be narrated.

Höllerer and Feiner (1997) have built a university campus information system, worn as a wearable computer. This device is endowed with a variety of sensors for head tracking and image registration. Both the size of the wearable, mounted on a large and heavy backpack, as well as the size of the display, are inappropriate use in a museum visit.

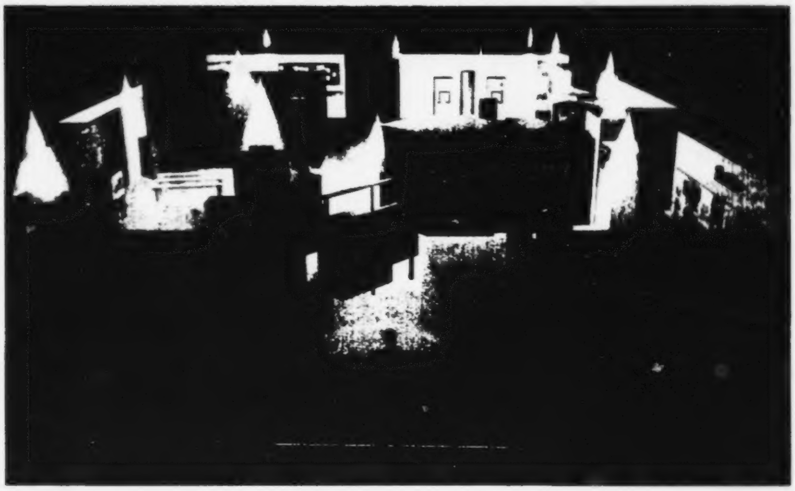

Figure 29. 3D model: start frame of animation of a visitor at the exhibit.

Various groups are working to augment the museum visit with mobile devices that are not wearable computers but handhelds, which do not have a private-eye or head-mounted display but rely on the handheld's screen for visual communication. One of the main drawbacks of such devices is that the visitor is obliged to toggle his/her attention between the objects on display and the handheld's screen, alternatively looking ahead towards the objects and then down to the screen. The private eye of the museum wearable instead allows the visitor to have a true augmented reality experience by presenting the viewer with a fused image which mixes together the real world and the computer augmentation either as a picture-in-picture effect or as two superimposed layers of information.

Amongst the handheld based projects, Spaspjevic and Kindberg (2001) describe the electronic guidebook, currently under development at the San Francisco Exploratorium Science Museum. This device uses a combination of infrared and RFID location sensors to give visitors the ability to either view or bookmark Web pages which provide additional description information on the objects of the exhibit. This portable device is used mainly to record the visitors' path through the exhibit, typically as a group, so that later the visit can be discussed and commented upon in a classroom setting.

The European HIPS project proposes a user-centered approach to information delivery in museums (Broadbent and Marti, 1997). As participants to HIPS, Oppermann and Specht (1999) describe an adaptive system which requires the user's intervention to personalize the presentation. This however can be quite disruptive of the visitor's museum experience, in a way similar to those multi-path DVD movies which stop at all turning key point and ask the viewer to choose an option before continuing. A system such as the museum wearable, capable of inferring the user's preferences with a mathematical model able to use the sensors' information as cues and the curators' knowledge of their public as a guideline, , can be indeed less disruptive.

The MUSEpad project (Kirk, 2001) is developing an original mobile device to enable visitors with disabilities to customize and optimize their learning and leisure experiences in museums. Their feasibility study examines user personalization to allow for adaptation to the needs of different users, whether by providing video-only, audio-only or magnified images for visitors with poor vision.

Aoki and Woodruff (2000) describe an electronic guidebook prototype that facilitates social interaction during the museum visit. Using a touch sensitive screen on the handheld and a photograph-based interface that the visitors can click on, the handheld offers a text or audio description of the selected object, according to the user's preferences. Visitors are then provided a mechanism to hear each other's audio selection. Such research on mobile augmen-

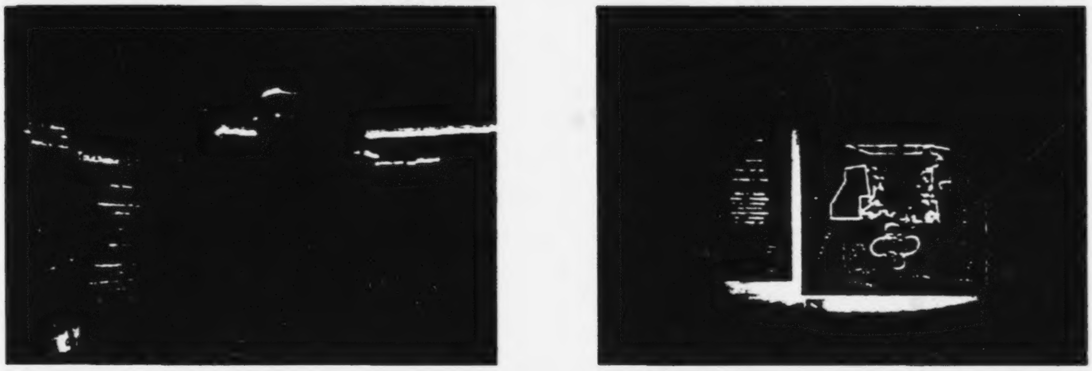

Figures 30, 31. Early version of the museum wearable at the SIGGRAPH 99 Millennium Motel.

tation devices which can be used collectively as a group is quite important for the future of this field.

The author showed an early prototype of the museum wearable called "Wearable City of News", as a demonstration for the SIGGRAPH 99 Millennium Motel, where hundreds of users tested the system for seven full days (http://www.siggraph.org/s99/conference/etech/projects.html). The system featured a jacket with an embedded computer and head-mounted display to show visitors Web pages and video previews of the Millennium Motel demonstrations using a short range but low power infrared location system [figures 30, 31].

8. Discussion and Future Work

The museum wearable fuses the audiovisual documentary, which illustrates and extends an exhibit, with the visitor's path inside that exhibit, using a wearable computer. By having the public use this device, curators and exhibit designers can accomplish multiple goals simultaneously: they can have objects narrate their own story; they do not needs special rooms to show audiovisual explanations about the exhibit as with the wearable the narrative is unfolded by the visitor's path in the museum; they can show more artwork than what is physically on display, including video, images, audio, and text about other important objects for the exhibit (these usually do not see the light because of the physical limitations of the available space); they do not need to disseminate panels with textual explanation or video monitors along the aisles of the exhibit as that information can now be tailored to each individual visitor; they can personalize the audiovisual explanations provided to the public based on the visitor's type and exploration strategy.

The museum wearable provides more than a simple associative coupling between inputs and outputs. The sensor inputs, coming from the long range indoors infrared positioning system, are coupled to digital media outputs via a user model, and estimated probabilistically by a Bayesian network. The ability to coordinate and present the visual material as a function of the visitor's estimated type (i.e. busy, greedy, or selective types, or other appropriate types), seamlessly, appropriately, and in conjunction with the path of the wearer inside the exhibit, is an important feature of this device. Bayesian networks have the additional advantage that they allow to us encapsulate our human knowledge about the context of use of the museum wearable (a particular exhibit, a trade show) in appropriate nodes of the network.

With respect to the traditional museum audio tour, the museum wearable introduces the following innovations: it does not constrain the visitor to follow any predefined sequential path in the museum, or to continuously press buttons, but it relies in its sensing system to find the visitor's location and respond; consequently, it adds a layer of visual augmentation, not just auditory; through Bayesian network based user modeling it provides personalized content to each visitor, as a function of the estimated visitor type.

65

An experimentation phase at the museum should follow the current research. A procedure must be set to establish if and how the museum wearable does actually enhance learning and entertainment at an exhibit, or how the content shown does actually match the visitor's preferences.

The research here presented can also be expanded in various ways. One direction of work is to find ways to get to know the visitors better so as to target content presentation more accurately towards their level of knowledge or competence. Asking visitors to fill out lengthy questionnaires upon entering the exhibit may not be practical. It is instead desirable to experiment with additional sensors, such as the GSR, a motion sensor (accelerometer), or a tiny video camera to obtain a more accurate estimate of the visitors' interest profiles and levels of attention. More visitor tracking data could be gathered at the museum site, to eventually infer more visitor types than the ones described in this document, and compare them with the more sophisticated visitor typologies discussed in the museum literature.

The museum wearable can also become a very useful tool to gather visitor tracking data in the museum. For example, rather than coming up with a set of visitor types from the museum literature, one could adopt the opposite approach of "inferring" the visitor types from a statistical analysis of the tracking data (path and stop duration) gathered by visitors with the museum wearable. This information would help the curator, exhibit designer, and the modeler of the interactive museum experience to refine their knowledge about visitor types for a specific exhibit. Similarly, by analyzing the posterior values of the object nodes, the curator and the exhibit designer could see which objects are the most interesting or boring for the visitors, and change the exhibit layout accordingly.

An important extension to the museum wearable would allow it to support visitors who want to come to the museum as a group and have the freedom to comment and discuss the artwork amongst themselves. A simple modification to the current prototype would be to add a small microphone capable of detecting when the visitor is talking, and then automatically pause the narration. If for example the group of visitors is composed of high school students, it would be useful, for learning purposes, to make each visitor profile available to the users at the end of the exhibit, and have the system regroup the visitors according to matching profiles, for further discussion. This same capability could also be made available to visitors in the museum's cafeteria at the end of their tour, to play matchmaking among those who wish to be involved. Alternatively the visitor's profile, path, and length of stay can be used to create a Web-based exhibit catalogue whose URL can be sent to the visitor as a personalized basis for further learning.

Acknowledgments

A group of dynamic MIT undergraduates contributed to the museum wearable at different stages of development. The author would like to thank Tracie Lee, Aman Loomba, Manuel Martinez, Eric Hilton, Chin Yan Wong, Audrey Roy, Anjali D'Oza, and Sarah Mendelowitz for the many hours of work they dedicated to this project. My deepest thanks also go to Janet Pickering, Beryl Rosenthal, and Janis Sacco from the MIT Museum, to help make this project a reality. Neil Gerschenfeld, Alex Pentland, Glorianna Davenport, Kent Larson, Walter Bender, and Ron McNeil have also supported this project.

References

Aoki, P.M. (2001) and Woodruff, A. The Conversational Role of Electronic Guidebooks. In: Proc. UBICOMP 2001, Springer-Verlag.

Broadbent, J. (1997), Marti, P. Location aware mobile interactive guides: usability issues. In: Proc. of International Cultural Heritage Informatics Meeting (ICHIM '97), Paris, France.

Dean, D. (1994). *Museum Exhibition: Theory and Practice*. London, Routledge.

Healey, J. (1999), Seger, J., and Picard R. Quantifying Driver Stress: Developing a System for Collecting and Processing Bio-Metric Signals in Natural Situations. In: Proceedings of the Rocky Mountain Bio-Engineering Symposium, April 16-18 1999.

Höllerer, T. (1999), Feiner, S., Pavlik, J. Situated Documentaries: Embedding Multimedia Presentations in the Real World, In: Proc. Third Int. Symp. on Wearable Computers (ISWC '99), San Francisco, CA, October 18-19, 1999, pp. 79-86.

Hooper-Greenhill E (1999). *Museums and their visitors*, London, Routledge.

Houbart, G. (1994). *Viewpoints on Demand: Tailoring the Presentation of Opinions in Video*. MIT Masters Thesis.

Jordan, M.I. (1999), editor. *Learning in Graphical Models*. The MIT Press.

Kirk J. (2001). Accessibility and New Technology in the Museum. In: D. Bearman and J. Trant (eds.) Museums and the Web 01 Proceedings, Seattle, USA.

Klein, L. (1986). *Exhibits: Planning and Design*. Madison Square Press, New York, pp70-71.

Oppermann, R. (1999), Specht, M. A Nomadic Information System for Adaptive Exhibition Guidance. In: D. Bearman and J. Trant (eds.): Proceedings of the International Conference on Hypermedia and Interactivity in Museums (ICHIM 99), Washington, September 23 – 25, pp. 103 – 109.

Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, San Mateo, CA.

Schiele, B. (1999), Oliver, N., Jebara, T. and Pentland, A. DyPERS: Dynamic Personal Enhanced Reality System. In Proceedings of Intl. Conference on Vision Systems (ICVS 99). Gran Canaria. Spain.

Serrell, B. (1996). The question of visitor styles. In: S. Bitgood (Ed.). Visitor Studies: Theory, Research, and Practice, Vol. 7.1. Jacksonville AL: Visitor Studies Association, pp. 48-53.

Smyth, P. (1997). Belief Networks, Hidden Markov Models, and Markov Random Fields: A Unifying View. Pattern Recognition Letters, vol. 18, number 11-13, pp. 1261-1268.

Sparacino F. (1999), Davenport G., Pentland A., Media Actors: Characters in Search of an Author. In: Proc. of IEEE Multimedia Systems '99, International Conference on Multimedia Computing and Systems (IEEE ICMCS'99), Centro Affari, Firenze, Italy 7-11 June 1999.

Sparacino, F. (2001). *Sto(ry)chastics: a Bayesian network architecture for combined user modeling, sensor fusion, and computational storytelling for interactive spaces*. MIT PhD Thesis.

Spaspjevic, M. (2001) and Kindberg, T. A Study of an Augmented Museum Experience. Hewlett Packard internal technical report. HPL-2001-178.

Starner, T. (1995), and Pentland, A. Visual Recognition of American Sign Language Using Hidden Markov Models. In: Proc. of International Workshop on Automatic Face and Gesture Recognition (IWAFGR 95). Zurich, Switzerland.

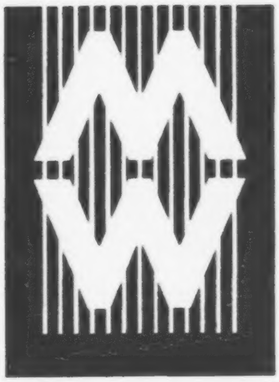

Foundations

Now That We've Found the 'Hidden Web,' What Can We Do With It? The Illinois Open Archives Initiative Metadata Harvesting Experience

**Timothy W. Cole, Joanne Kaczmarek, Paul F. Marty,
ChristopherJ. Prom, Beth Sandore, and Sarah Shreeves,
University of Illinois at Urbana-Champaign, USA**

Abstract

The Open Archives Initiative (OAI) Protocol for Metadata Harvesting (PMH) is designed to facilitate discovery of the "hidden web" of scholarly information such as that contained in databases, finding aids, and XML documents. OAI-PMH supports standardized exchange of metadata describing items in disparate collections such as those held by museums and libraries. This paper describes recent work done by the University of Illinois Library, recipient of one of seven OAI-related grants from the Andrew W. Mellon Foundation. An overview is given of the process used to export metadata records describing holdings of the Spurlock Museum at the University of Illinois. These metadata records were initially created to help track artifacts as they were procured, stored, and displayed and now are used also to support end-user searching via the Spurlock Museum Website. Spurlock metadata records were mapped to Dublin Core (DC) and then harvested into the Illinois project's metadata repository. The details of the processes used to transform the Spurlock records into OAI compliant metadata and the lessons learned during this process are illustrative of the work necessary to make museum collections available using OAI-PMH. Assuming institutions like Spurlock make metadata available, what then can be done with these information resources? We discuss the OAI-based search and discovery services being developed by the University of Illinois. Issues such as need for normalization of metadata, importance of presenting search results in context, and difficulties caused by institution-to-institution variations in metadata authoring practices are discussed.

Keywords: Open Archives Initiative (OAI), metadata harvesting, Dublin Core, cultural heritage, interoperability, heterogeneous collections

Introduction

Although the Open Archives Initiative (OAI) originally focused on the exchange of metadata describing e-prints in the scientific community (Van de Sompel & Lagoze, 2000), the OAI-Protocol for Metadata Harvesting (OAI-PMH) holds much promise for similar exchanges of metadata describing the collections of museums and other cultural heritage institutions (Perkins, 2001). Materials in these types of collections often are not well indexed (or not indexed at all) by commercial Web search engines. Metadata describing such holdings is hidden in databases, finding aids, and XML documents, or otherwise is not readily available to Web search systems, which typically understand little more than simple HTML. As a result, these materials remain hard to find and out of reach for many researchers. With an interest in making these materials more visible to scholars and other researchers, the University of Illinois at Urbana-Champaign, through an Andrew W. Mellon Foundation grant, is exploring the feasibility of using OAI-PMH to build services to reveal and make more accessible collections of cultural heritage material.

While a focus of the Illinois research is on building and testing software tools designed to harvest metadata provided by OAI-compliant metadata providers, we also are exploring what can be done with OAI-harvested metadata. What is the fidelity and usefulness of the indexes and search services that can be built atop a repository of metadata harvested using the OAI-PMH? The utility of such a metadata repository will depend on many factors, including the development of more uniform metadata authoring practices across various communities of metadata providers. Metadata is currently used primarily for a variety of local purposes, and the schemas to which collections conform have been adapted and tailored to meet local needs. Less attention has been paid to using metadata to support

BEST COPY AVAILABLE

Museums and the Web 2002: Proceedings
Archives & Museum Informatics, 2002, p. 63

universal interoperability. With the advent of OAI and similar protocols, mapping local schemas to more universally standard schemas for the support of interoperability is gaining ground.

Even assuming broader and more consistent application of metadata schemas designed to support interoperability, there remains considerable work to be done in learning how best to utilize aggregations of metadata describing heterogeneous information resources. There is a need to normalize metadata so as to enable more consistent retrieval of results from cross-collection searches. Challenges also exist for presenting search results in ways that provide appropriate contextual information for the records retrieved. We will discuss more detailed examples of these and other issues throughout the paper. We will outline our efforts to overcome these issues in building a useful cross-collection repository of cultural heritage materials. We will also discuss our work in progress, our plans for future development, and our reasons for believing that the OAI protocol has great potential for increasing access to and exposure of hidden resources via the Web.

Illinois Project Harvest Experience To Date

As of February, 2002, the Illinois OAI-PMH project has harvested metadata from twenty-five different institutions or consortiums. The resources described range from museum artifacts such as pottery and clothing, to archival manuscript and personal paper collections, to digitized photographs and monographs. The sets of harvested metadata range in size from the largest at 900,000 records to the smallest at 23 records. Some sites include metadata not relevant to cultural heritage mixed with more appropriate metadata. Our OAI Harvester software excluded such non-relevant metadata from our indices. In total, we have inspected over two million individual records, resulting in an index of approximately 770,800 records relevant to the subject domain of interest. The institutions providing metadata are diverse, including academic libraries, museums, historical societies, public libraries, the Library of Congress, and special consortia such as the Colorado Digitization Project and the Alliance and Lincoln Trail Library systems in Illinois. Eleven contributing institutions are registered OAI metadata providers <http://www.openarchives.org/Register/BrowseSites.pl>. Fourteen institutions have provided the Illinois OAI-PMH project team with snapshots of relevant metadata, in a few cases providing an entire database, which has then been made available from Illinois servers in a manner compliant with OAI-PMH.

While all of the metadata harvested from the registered OAI metadata provider sites was in the required Simple Dublin Core format (DC), metadata from the fourteen unregistered sites were given to the project in a variety of formats. These included finding aids formatted in encoded archival description (EAD), MARC format metadata describing bibliographic information, and local metadata schemas stored in HTML files, XML files, or proprietary database structures. We subsequently mapped these records into DC, prior to making them available in accord with the OAI-PMH.

Cross-Collection Repository Issues

To enhance cross-collection searching, local metadata schemas need to be mapped to standard schemas. Even with greater adoption of standard metadata schemas like DC, there remain wide variations in the use of authoring conventions and the depth of descriptive information included when creating metadata to describe cultural heritage collections.

OAI-PMH Basics

The Open Archives Initiative Protocol for Metadata Harvesting (OAI-PMH) is a new and still evolving protocol designed to allow institutions to share metadata easily. It is hoped that the standard can help increase the discoverability of resources by scholarly researchers. The Protocol underwent one minor revision during its first full year of experimental implementation. Version 2.0 is scheduled for release in May 2002. The fundamental pieces of the protocol are its adherence to well-formed XML, its use of the HTTP standards for data transmission, and its requirement that metadata records shared via the OAI-PMH be made available in the DC metadata schema (optionally records may be made available in additional metadata schemas as well) (Lagoze and Van de Sompel, 2001).

	LIBRARY OF CONGRESS		AMERICAN MUSEUM OF NATURAL HISTORY		CIMI DEMO REPOSITORY		SPURLOCK MUSEUM	
Sample Analyzed	**79,629 records**		**2,004 records**		**197,233 records**		**46,612 records**	
	% of records containing element	*Average times used per record*	*% of records containing element*	*Average times used per record*	*% of records containing element*	*Average times used per record*	*% of records containing element*	*Average times used per record*
Title	100%	1.00	100%	1.00	71%	1.07	100%	1.00
Type	100%	2.34	100%	1.00	100%	3.50	100%	3.00
Creator	97%	1.17	100%	1.00	43%	1.03	16%	1.00
Publisher	73%	1.00	100%	1.00	100%	1.37	100%	1.00
Description	23%	2.39	4%	1.00	94%	1.13	93%	3.44
Source	Not Used	-	Not used	-	25%	1.00	2%	1.00
Relation	Not Used	-	100%	1.00	96%	1.45	23%	1.03
Format	Not Used	-	100%	1.00	22%	2.54	70%	1.00
Date	72%	1.00	100%	1.00	64%	1.10	70%	1.00
Subject	66%	1.61	48%	1.00	95%	3.21	100%	1.00
Identifier	100%	1.00	100%	1.00	100%	1.43	100%	2.00
Language	56%	1.00	Not Used	-	60%	1.00	Not Used	-
Coverage	40%	1.07	100%	1.00	60%	3.70	99%	3.23
Rights	25%	1.00	100%	1.00	81%	1.65	100%	1.00
Contributor	Not Used	-	100%	1.00	44%	2.19	50%	1.00

Table 1 - Dublin Core elements used and number of times repeated

One of the common challenges in setting up an OAI-PMH metadata provider service is ensuring that the metadata to be exported is mapped properly into DC. There are fifteen DC elements (version 1.1): Title, Creator, Subject, Description, Publisher, Contributor, Date, Type, Format, Identifier, Source, Language, Relation, Coverage, and Rights. Definitions and recommended usage of these elements can be found at the Dublin Core Metadata Initiative Website <http://www.dublincore.org>. Any or all of these elements may be used, and all are repeatable. Because much of the existing metadata about museum and special collections was developed with limited resources, using local schemas and with local needs in mind, the process of mapping to DC can be time and resource consuming and potentially frustrating. There is, however, a growing community of DC users and freely available tools to assist with this process.

Metadata Authoring Practices

Harvesting done so far shows that metadata authoring practices differ both in the selection of DC elements used and the depth of information provided even within specific communities. As previously noted, DC elements are both optional and repeatable. These allow a significant degree of variation in the interpretation of which elements to use and how to use them when metadata are created in DC and when mapped to DC. We examined records from the Library of Congress American Memory Project, the American Museum of Natural History, the CIMI Demonstration Repository, and the Spurlock Museum to compare the use of DC elements (See Table 1). Within each set of sample records, we calculated the percent of records that contain each of the DC elements and the average number of times an element is used in records that contain it. The usage variations are due in part to

BEST COPY AVAILABLE

the intrinsic differences in the collections described by the metadata. However, some variations are clearly due to different decisions made by each institution when determining how to map their metadata into DC. These variations pose a challenge when attempting to build cross-collection search services and determining what depth of detail should be shown in record display. The variations can also influence how a search engine's ranking algorithm orders the result set.

Spurlock Museum Metadata Mapping

Museums began to use personal computers for tracking information about their collections using a variety of proprietary databases. The use of proprietary software not specifically designed with the museum community in mind can lead to inconsistencies in the way collections are catalogued and tracked. The descriptive elements within a locally created metadata format and the degree of completeness when classifying and cataloging the items can vary widely from one organization to another. This may have the effect of diluting the potential for discoverability as searching aggregated repositories requires some degree of predictable structure and normalization of metadata terms or concepts. At the same time, one should not overlook the value of descriptive, locally created metadata applied to artifacts by professionals familiar with the collection and the museum. This is particularly valuable for collections used primarily by a local community of users. At the University of Illinois Spurlock Museum, FileMakerPro software is used to track the procurement, storage, stage of processing, and display of artifacts. It is also used to maintain descriptive information about the materials. The FileMakerPro database interacts with a Web server to provide public access to a portion of the collection <http://www.spurlock.uiuc.edu/>. This database provides direct online access to 45,000 of over 210,000 records describing cultural heritage and natural science artifacts. The Illinois OAI-PMH team was provided with metadata for all 210,000 records.

As a first step to making these metadata records available via OAI, Spurlock's metadata were extracted from the FileMakerPro database using a simple extraction script. It was necessary to map the locally created and customized metadata format to a schema more closely related to DC. To preserve richness of the local metadata schema, the Spurlock metadata had to be mapped first to Qualified Dublin Core (DCQ) and then to simple DC (Dublin Core Qualifiers, 2000).

The decisions made about mapping the metadata were more time consuming than the time spent writing the scripts and manipulating the data. Here is an example of some of the decisions that were made for the cultural materials provided by Spurlock:

- **Subject**

One instance of the DC Subject element was included in each record. The decision was made to concatenate three fields from the original metadata records that could be considered equivalent to DC Subject. Colons were used to distinguish each of the three original strings that were concatenated to create the single DC Subject element. If there were not three values to concatenate, then the DC Subject filed would have two colons next to each other or a lone colon at the beginning or end of the element.

- **Date**

The Date element was qualified as DCQ date.created.

- **Coverage**

The Coverage element was qualified by appending Spurlock metadata field names to the DC element as DCQ coverage.spatial, coverage.temporal, or coverage.cultural.

- **Description**

Key words (Materials, Manufactural Process, Munsell Color ID) were added to the Description field values in the original metadata to clarify what was meant by the content of the Description element.

Table 2 gives the complete list of mappings from local Spurlock-specific metadata schema to DC metadata schema used to export the metadata via the OAI- PMH. In addition to the mapping given below, 3 fixed DC Type values ("cultural," "physical object," and "original") were included in each Spurlock metadata record exported via OAI. Figure 1 shows how a typical Spurlock metadata record looks when retrieved by an end-user using the Spurlock Museum website. Figure 2 shows this same metadata record exported via OAI.

72

SPURLOCK DATABASE FIELD	DC ELEMENT
Object Name	Title
Artist/Maker	Creator
Classification	Subject
Measurements	Format
Culture	Coverage
Geographic	Coverage
Time Period	Coverage
Date	Date
Materials	Description
Manufacturing Process	Description
Munsell Color ID	Description
Visual Description	Description
Published Description	Description
Scholarly Notes	Description
Reproduction (if it is a reproduction)	Description
Reproduction Remarks (if it is a reproduction)	Source
Bibliography	Relation
Comperanda	Relation
Credit Line/Donor Name (if publicly assessable)	Contributor
Spurlock Museum, University of Illinois at Urbana-Champaign	Rights
Spurlock Museum, University of Illinois at Urbana-Champaign	Publisher
Accession Number	Identifier
URL to Spurlock's Database	Identifier

Table 2 - Spurlock cultural material metadata fields mapped to Dublin Core

Item-Level vs. Collection-Level Metadata

One of the more challenging aspects of implementing the OAI protocol is mapping from metadata schemas designed to describe collections of materials (e.g., an EAD Finding Aid record) to the DC metadata schema. Finding aids may describe as many as several thousand items or folders in an archive while DC has typically been used to describe individual items (e.g., books, photographs, letters, personal journals, audio files). Each EAD record includes metadata describing the entire collection and a "description of subordinate components" which lists the separate series, sub-series, folders and items found in the collection. Some EAD files reach hundreds of kilobytes, or even several megabytes, in size. The challenge is to allow the richness of such a large file to be exposed and made searchable alongside other records that describe a single item or a much smaller collection.

Another challenge is due to the very flexible nature of the EAD standard designed to encourage participation by a large number of institutions. It is not surprising that the EAD records analyzed for the Illinois OAI-PMH project reveal many differences in tag structures and encoding patterns between institutions. Despite these differences, the records seem to have enough consistency to make a general cross mapping between EAD and OAI-DC possible. Although the EAD Application Guidelines include two recommended mappings to DC, one for the finding aid itself (i.e. the electronic resource) and another for the resources described in the finding aid (i.e. the actual manuscripts or archives), neither mapping will generate records adequate for use in an OAI repository. While researchers are interested in discovering the existence of actual archives or manuscripts, metadata about the finding aid itself is also necessary in order to point effectively to the source records. For this reason, a flexible approach in mapping an EAD finding aid into an OAI record is needed.

We are currently testing a schema that produces many DC metadata records from a single EAD file. We first produce a record describing the entire collection of materials in the finding aid. This top-level record becomes a base record for the finding aid and must be as complete, accurate, and concise as possible. The base record includes a link to the source EAD file as well as references to related parts of the collection. These related parts are described in other individual records we produce for each component level found within the EAD file's description of subordinate components. Although it is not possible to link these records to each other adequately using simple DC, a sophisticated use of qualified DC fields can produce linked records. We hope this method will provide functional and easily searchable records.

Enhancing Discoverability of Resources

Metadata authoring practices as discussed above play important roles in determining the value of a cross-collection repository focusing on cultural heritage materials. Providing context for the metadata, normalization techniques, and developing the search engine and interface also need to be examined and explored for ways to enhance the discoverability of the collected metadata.

BEST COPY AVAILABLE

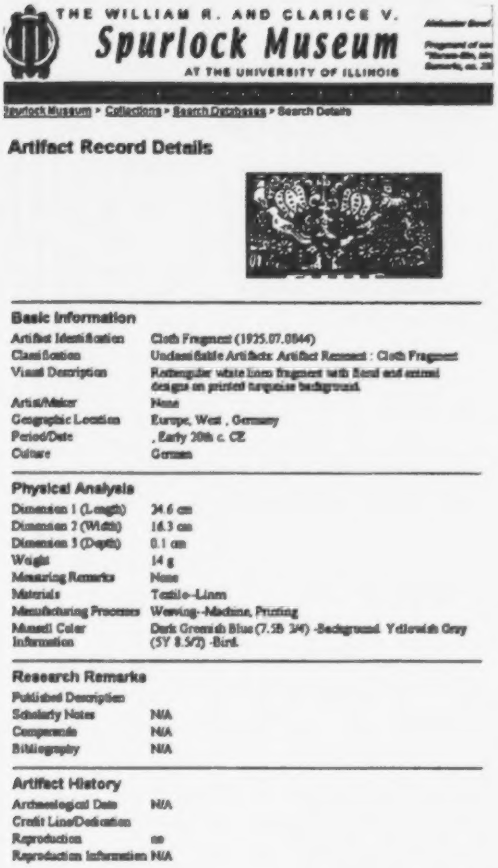

THE WILLIAM R. AND CLARICE V.
Spurlock Museum
AT THE UNIVERSITY OF ILLINOIS
Spurlock Museum > Collections > Search Databases > Search Details

Artifact Record Details

Basic Information

Artifact Identification	Cloth Fragment (1925.07.0044)
Classification	Unclassifiable Artifacts: Artifact Remnant : Cloth Fragment
Visual Description	Rectangular white linen fragment with floral and animal designs on printed turquoise background.
Artist/Maker	None
Geographic Location	Europe, West , Germany
Period/Date	, Early 20th c. CE
Culture	German

Physical Analysis

Dimension 1 (Length)	24.6 cm
Dimension 2 (Width)	16.3 cm
Dimension 3 (Depth)	0.1 cm
Weight	14 g
Measuring Remarks	None
Materials	Textile--Linen
Manufacturing Processes	Weaving--Machine, Printing
Munsell Color Information	Dark Greenish Blue (7.5B 3/4) -Background. Yellowish Gray (5Y 8.5/2) -Bird.

Research Remarks

Published Description	
Scholarly Notes	N/A
Comparanda	N/A
Bibliography	N/A

Artifact History

Archaeological Data	N/A
Credit Line/Dedication	
Reproduction	no
Reproduction Information	N/A

Figure 1 - Record shown in Spurlock database

Providing Context for Content

One real concern related to the development of cross-collection repository search tools is the context of the original artifact or digital resource. The value of information about the materials that may be found in museum collections and archives is often related to the context and the provenance of the object(s). For example, a photographic slide depicting part of a remaining wall of a basilica dating to the 4th Century would have less value if it were viewed out of the context of the architectural structure it represents. It is not good enough simply to make the slide available for viewing; it needs to be put into a proper context. Ideally, the holding institution of the artifact or digital resource will determine the appropriate context. We attempt to maintain context in our cultural heritage repository by providing external links back to the holding institutions. When our repository encounters URLs embedded within any of the metadata fields, they are mapped to the DC Identifier element. The links sometimes lead to the holding institution or organization's Web site, to the digital display of the resource's record from the holding institution or organization's database, or to an actual display of the resource itself. Given the non-persistent nature of many URLs, it is possible that providing access to them may lead to a large number of dead links over time. With regularly scheduled re-harvesting of the metadata, changes to the URLs should be reflected in the records and therefore allow us to avoid excessive dead links.

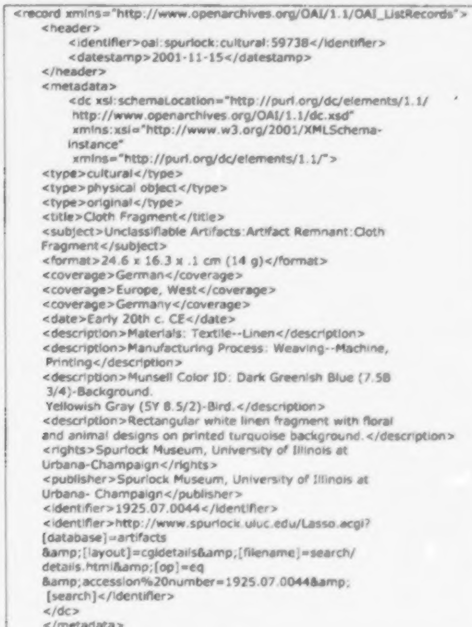

```
<record xmlns="http://www.openarchives.org/OAI/1.1/OAI_ListRecords">
    <header>
        <identifier>oai:spurlock:cultural:59738</identifier>
        <datestamp>2001-11-15</datestamp>
    </header>
    <metadata>
        <dc xsi:schemaLocation="http://purl.org/dc/elements/1.1/
         http://www.openarchives.org/OAI/1.1/dc.xsd"
         xmlns:xsi="http://www.w3.org/2001/XMLSchema-
        instance"
         xmlns="http://purl.org/dc/elements/1.1/">
    <type>cultural</type>
    <type>physical object</type>
    <type>original</type>
    <title>Cloth Fragment</title>
    <subject>Unclassifiable Artifacts:Artifact Remnant:Cloth
    Fragment</subject>
    <format>24.6 x 16.3 x .1 cm (14 g)</format>
    <coverage>German</coverage>
    <coverage>Europe, West</coverage>
    <coverage>Germany</coverage>
    <date>Early 20th c. CE</date>
    <description>Materials: Textile--Linen</description>
    <description>Manufacturing Process: Weaving--Machine,
     Printing</description>
    <description>Munsell Color ID: Dark Greenish Blue (7.5B
     3/4)-Background.
     Yellowish Gray (5Y 8.5/2)-Bird.</description>
    <description>Rectangular white linen fragment with floral
    and animal designs on printed turquoise background.</description>
    <rights>Spurlock Museum, University of Illinois at
    Urbana-Champaign</rights>
    <publisher>Spurlock Museum, University of Illinois at
    Urbana- Champaign</publisher>
    <identifier>1925.07.0044</identifier>
    <identifier>http://www.spurlock.uiuc.edu/Lasso.acgi?
    [database]=artifacts
    &amp;[layout]=cgidetails&amp;[filename]=search/
    details.html&amp;[op]=eq
    &amp;accession%20number=1925.07.0044&amp;
     [search]</identifier>
    </dc>
    </metadata>
</record>
```

Figure 2 - Record shown as exported by OAI-PMH

For an example of a repository that has made thoughtful consideration of these efforts, consider the metadata collection contributed to the Illinois OAI-PMH project by TDC, the Teaching with Digital Content project <http://images.library.uiuc.edu/projects/tdc/>. This project provides a searchable database of images that can be used by K-12 teachers to develop teaching modules that meet specific curriculum requirements. Partners providing metadata for the project include museums, the Illinois State Library, the Chicago Public Library and others. All the metadata has been entered into a database in a format that meets the specific needs of the TDC project and does not conform to DC. However, the display of the records is very clean and consistent and maps well into DC for the purposes of the Illinois OAI-PMH project. With a link directly to the Teaching with Digital Content resource provided from a URL that has been mapped into the DC Identifier field, the context is easily maintained by providing more background information about the record when the user clicks on this link.

We plan to explore methods for providing context internally within the repository as well, linking back to the owning institutions. Providing easy access to related records may reveal overlooked information and insights. The collocation of works by the same author and under the same subject headings is a given in databases using a standardized metadata format and data entry. In an aggregated database, this collocation task is not trivial, but is potentially even more illuminating. The context could be thought of in different ways, such as:

- a specific archaeological dig;
- a specific time period or era (e.g. Civil War);
- a single donor;
- a particular geographic area (e.g. South America); or
- a specific genre of art or literature (e.g. Art Nouveau).

Metadata Normalization

In order to provide context internally and to enhance discoverability of metadata records in a cross-collection repository, some normalization of the metadata is desirable. We believe that normalization can be done on several of the DC elements, including Type, Format, Coverage, and Date. We hope to provide some degree of normalization on the Subject element but have not yet developed our strategy for this more complex normalization.

Effective normalization requires us to:

- Understand how the element was interpreted by metadata providers and which elements in other metadata formats were mapped to the DC element;
- Identify which – if any – vocabularies were used by data providers;
- Determine whether there is any controlled vocabulary that the project team could successfully apply to all of the data providers, or, if not, create such a vocabulary specific to our repository;
- Apply our vocabulary to the metadata to augment the 'native' vocabulary;
- Build mechanisms into the search interface that would take advantage of the normalization; and
- Gauge the success of the normalization for resource discovery.

These goals translate into a five-step process as follows.

1. Extract and analyze the element values (content)

We organized the element values by metadata provider. Each unique value was listed, along with the number of times it appeared within each individual metadata set. For example, we discovered that only eleven of the twenty-five metadata providers used the Type element and that approximately 1440 different Type values appear in the entire metadata set. Some providers only use one value, and others use over 800. (See Table 3) The number of values seems to be dependent on whether the institutions are using a specific or general vocabulary. A large number of records in the CIMI metadata use very specific types such as "Physical Object:TOYS."

2. Determine how each element is interpreted and what controlled vocabulary, if any, is used.

We discovered that metadata providers used elements in a variety of ways. For example, the DC Date element was used for the date the digital item was created, the date the physical item was created or published, and the date an item was added to a collection. Because OAI-PMH requires use of simple DC, the qualifiers that may have explained these values further were not present. Similarly the variety and specificity or generality of controlled vocabularies also influenced the next steps in the normalization process.

METADATA PROVIDER	NUMBER OF VALUES	VOCABULARY USED FOR TYPE ELEMENT AND SPECIFICITY OF VOCABULARY
Alliance Library System	431	LC Subject Headings and ALA's Guidelines on Subject Access to Individual Works of Fiction, Drama, etc. (specific)
American Museum of Natural History	1	DCT1 (general)
Celebration of Women Writers	1	DCT1 (general)
CIMI Demonstration Repository	886	CIMI modifications to Dublin Core Type Vocabulary (DCT1) and local vocabularies (general and specific)
Formations	1	Local vocabulary (general)
Library of Congress American Memory	87	LC Thesaurus of Graphic Materials II (specific)
Open Video Project, University of North Carolina – Chapel Hill	1	Local vocabulary (general)
Perseus Digital Library	1	DCT1 (general)
Spurlock Museum	4	CIMI modifications to DCT1 (general)
University of Michigan, Digital Library Production Service	1	Local vocabulary (general)
University of Tennessee Special Collections	25	Local vocabulary (hierarchical)

Table 3 - Number of Values and Vocabulary Used for Type Element

3. Determine focus and vocabulary for normalization

Each DC element was examined to determine our focus for normalization for that particular element. While the Date element suggests a need to choose between options like 'date created' or 'date contributed' and the particular format the value is displayed in, the Type element required us to focus on the range and specificity of vocabulary used. For the Type element we determined resource discovery may be enhanced by adding slightly more general terms into the record alongside the 'native' vocabulary (e.g.: 'physical object' added to 'toy'). Following the guidelines from the DC Type Working Group, we developed our own vocabulary, pulling terms from existing vocabularies where applicable, that provided the level of specificity that would benefit our users the most (Apps, 2001).

4. Normalizing the data

Once the vocabulary is developed, the values in the metadata are mapped to the vocabulary terms. The actual normalization of the data is an automated process. Following the mapping, values from the vocabulary are added to each record as additional. After new records are harvested, the system compares the values of specific elements with the mapping already done and adds appropriate terms from our vocabulary where applicable. If a term that has not already been mapped appears, it is flagged and looked at by human eyes.

5. Providing services based on the normalization process

The last step is to provide services so that users can take advantage of the normalization process. The ability to limit searches or sort results by type of resource or to search for specific date ranges or geographic areas is each made possible by the normalization process.

End-User Search Interface Design Strategies

We began our interface design following guidelines for scenario-based design (Carroll, 2000). We identified likely users of the system as two distinctly separate groups - scholarly researchers and K-12 teachers and students. We have prioritized the scholarly researcher as our primary user group for the preliminary interface design phase of the project.

BEST COPY AVAILABLE

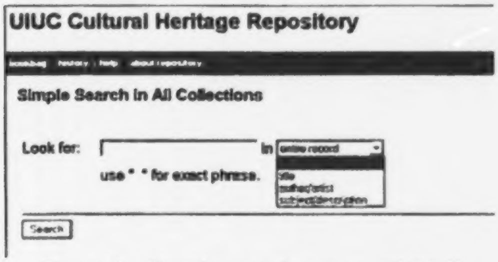
UIUC Cultural Heritage Repository
bookbag | history | help | about repository
Simple Search in All Collections
Look for:
in
entire record
title
author/artist
subject/description
use " " for exact phrase.
Search

Figure 3 - Simple search screen <http://oai.grainger.uiuc.edu/oai/search>

As time allows, we also intend to focus on K-12 teachers and students. In both cases, the need for a simple design seems important. This is in line with generally accepted "best practices" in interface design and is a logical choice given the potential for complex and/or inconsistent record displays due to the heterogeneity of the materials and institutions represented in the repository. We identified several types of common tasks our users might want to do with a repository of cultural heritage materials and sketched out scenarios that supported these tasks. Based on these scenarios and working from the interface built into our search engine (developed by the University of Michigan's Digital Library Extension Service), we designed the first iteration of the end-user search interface following general usability heuristics as outlined by Nielsen (2000). Figure 3 shows a simple, preliminary search interface developed as much for diagnostic purposes as for the end-user.

Work in Progress / Future Plans

As of February 2002, we have begun usability testing. The feedback provided has already highlighted obvious flaws both in the actual interface as well as back-end indexing methods. We will continue usability testing with new iterations of the interface. Subjects will include students enrolled at the University of Illinois Graduate School of Library and Information Science. We will also participate in a survey of potential users to be distributed jointly by University of Illinois and the University of Michigan. This survey is intended for various scholarly researchers on both campuses and will provide feedback to be used in the future design and enhancement of the system. Current plans for enhancing search functionality include adding options to combine and refine searches and access to stored "expert searches." These expert searches are intended to provide examples of search strategies used by experts from specific user communities. We will also include in our development user search features that focus on the K-12 teacher and student user profiles as prioritized by our original scenario-based design decisions. This user group has particular needs that center on curriculum development and the use of on-line resources for achieving clearly defined, institutionally driven educational outcomes. Under consideration is an enhancement option that would gather user input in the form of annotations or other types of notes one might be inclined to make about particular sets of retrieved records and add these to the index in a way that provides links and displays relationships between records that are related according to the researchers using the system.

Conclusion

The Illinois OAI Metadata Harvesting Project represents an attempt to enable the integrated searching of diverse types of cultural heritage information. In so doing, scholars may discover new links across materials that are physically dispersed and that may potentially address common themes across disciplines that have not been previously formally recognized. Providing cross-collection repositories that preserve the context of the items represented by disparate metadata and offer easily navigable search interfaces of this metadata requires several essential components:

- Metadata authoring practices must be compliant with community accepted standard schema(s);
- Techniques for displaying item-level vs. collection-level records need to be developed;
- Normalization processes should be applied to the metadata prior to indexing; and
- 'Best practices' for interface design should be adhered to for end-user search tools.

The value of a usable search system that represents texts, manuscripts, images, digital objects, and artifacts simultaneously is difficult to quantify at this stage of our investigation. However, we believe that the exploration of techniques for providing simultaneous access to such a range of materials is essential if we are to assist today's scholars in achieving their full potential as researchers.

BEST COPY AVAILABLE

Acknowledgement

The authors would like to acknowledge the generous support of the Mellon Foundation and the willingness of numerous institutions to provide metadata allowing us to explore the potential of OAI-PMH for cultural heritage materials.

References

Apps, Ann, editor (2001). Guidance for Domains and Organizations Developing Vocabularies for Use with Dublin Core: Outline. Working Draft. 28 Nov 2001. Accessed February 4, 2002, from http://epub.mimas.ac.uk/DC/typeguide.html.

Carroll, J.M. (2000). Five reasons for scenario-based design. *Interacting with Computers* 13: 43-60.

Dublin Core Metadata Initiative. http://www.dublincore.org/.

Dublin Core Element Set, Version 1.1: Reference Description. 1999-07-02. http://www.dublincore.org/documents/dces/.

Dublin Core Qualifiers. 2000-07-11. http://www.dublincore.org/documents/dcmes-qualifiers/.

Encoded Archival Document Application Guidelines for Version 1.0. 1999. http://lcweb.loc.gov/ead/ag/aghome.html.

Lagoze, C. & Van de Sompel, H. (2001). The open archives initiative: building a low-barrier interoperability framework. *Proceedings of the first ACM/IEEE-CS joint conference on Digital libraries 2001*. Retrieved February 8, 2002 from ACM Digital Library. http://doi.acm.org/10.1145/379437.379449.

Nielsen, J. (2000). Designing web usability: the practice of simplicity. Indianapolis, IN: New Riders Publishing.

Open Archives Initiative Protocol for Metadata Harvesting. Version 1.1. 2001-07-02. http://www.openarchives.org/OAI/openarchivesprotocol.htm.

Open Archives Initiative Registered Data Providers. http://www.openarchives.org/Register/BrowseSites.pl

Spurlock Museum at the University of Illinois. http://www.spurlock.uiuc.edu/.

University of Illinois at Urbana-Champaign OAI Metadata Harvesting Project. http://oai.grainger.uiuc.edu/.

Van de Sompel, H & Lagoze, C. (2000). The Santa Fe Convention of the Open Archives Initiative. *D-Lib Magazine* 6(2). http://www.dlib.org/dlib/february00/vandesompel-oai/02vandesompel-oai.html

Combining the CIDOC CRM and MPEG-7 to Describe Multimedia in Museums

Jane Hunter, DSTC Pty Ltd, University of Queensland, Australia

Abstract

This paper describes a proposal for an interoperable metadata model, based on international standards, which has been designed to enable the description, exchange and sharing of multimedia resources both within and between cultural institutions. Domain-specific ontologies have been developed by two different ISO Working Groups to standardize the semantics associated with the description of museum objects (CIDOC Conceptual Reference Model) and the description of multimedia content (MPEG-7) - but no single ontology or metadata model exists for describing museum multimedia content. This paper describes an approach which combines the domain-specific aspects of MPEG-7 and CIDOC-CRM models into a single ontology for describing and managing multimedia in museums. The result is an extensible model which could lead to a common search interface and the open exchange, sharing and integration of heterogeneous multimedia resources distributed across cultural institutions.

Keywords: Multimedia, Metadata, Interoperability, MPEG-7, CIDOC-CRM

1. Introduction

Multimedia provides museums with a rich paradigm for capturing, communicating and preserving cultural information. It offers new capabilities for structuring, interpreting and communicating knowledge, and the significance of artifacts within museum collections through the use of digital video, audio, images, graphics and animation. Making collections available in digital form, both in-house and through networks, provides museums with a tremendous opportunity to meet their educational mandate. When linked together over networks, museum multimedia databases become even more valuable as cross-cultural resources for educational and research purposes.

In addition, the potential to re-use multimedia content to create new intellectual property, has further accelerated the growth in the size and number of institutional multimedia databases. Existing multimedia objects are being combined and reused to generate complex, interactive multimedia, hypermedia, virtual reality displays and participatory exhibitions. This has led to a demand for systems and tools which can satisfy the more sophisticated requirements for storing, managing, searching, accessing, retrieving, sharing and tracking complex multimedia resources.

Metadata is the value-added information which documents the administrative, descriptive, preservation, technical and usage history and characteristics associated with resources. It provides the underlying foundation upon which digital asset management systems rely to provide fast, precise access to relevant resources across networks and between organisations. The metadata associated with multimedia objects is infinitely more complex than simple metadata for resource discovery of simple atomic textual documents and the problems and costs asssociated with generating such metadata are correspondingly magnified.

Metadata standards enable interoperability between systems and organizations so that information can be exchanged and shared. Standardized metadata models have been developed to describe museum objects (CIDOC Conceptual Reference Model) (CIDOC Conceptual) and to describe multimedia content (MPEG-7) (ICS) but no standards currently exist for specifically describing museum multimedia content. Hence the key goal of this project is to analyze and evaluate each of these existing standards and to determine a way to merge the two ontologies to generate a standardized model for describing museum multimedia content. Such a model, which is capable of supporting the exchange of information between existing collection management systems (for physical artefacts) and emerging digital asset management systems would enable knowledge and resources to be shared, re-used and exchanged to a much greater extent than is currently possible both within and between museums.

BEST COPY AVAILABLE

Hence in the remainder of this paper we describe both the CIDOC/CRM and MPEG-7 metadata models. We then analyze them both to determine the overlaps, intersections and differences. Based on this analysis we hope to be able to determine how the models can best be merged to combine the two domain specific vocabularies of MPEG-7 and CIDOC/CRM, without introducing semantic inconsistencies or redundancies. The final outcome is in essence, a single, machine-understandable, extensible ontology designed to support the description and management of multimedia resources within museums.

2. The Nature of Multimedia in Museums

2.1 Types of Multimedia in Museums

Audiovisual or "multimedia" content within museums is highly diverse and varies widely in origin, genre, purpose, media type, format, quality, age, context and the reason for its cultural significance or retention within a museum or collection. Multimedia in museums can include everything from disintegrating maps on paper to full feature films on DVD. Within the scope of this paper we are referring to: images, audio, video, multimedia, graphics and animation - in both analog and digital form. Table 1 below provides an overview of the typical constitution of multimedia collections within a museum.

In most cases, it is the semantic content, which is depicted or recorded on the multimedia resource which is of value i.e., the resource records an event, place, person, object or concept which is of cultural, historical, geographic or educational importance. For some multimedia resources, the cultural or historical significance lies with the medium or recording technology e.g., the earliest photographs, film, video and audio recordings. In some cases, the multimedia object is valuable because of the person who captured or recorded it (e.g., home movie collection of John F. Kennedy) or because it is part of a larger collection or bequest or it is exemplary of a particular genre, era or technique. In many cases, both the semantic content and the context, agents, technique and medium are all of importance and need to be recorded.

Images	photographs, prints, maps, manuscripts, documents, drawings, paintings, movie stills, posters
Audio	songs, music, plays, interviews, oral histories, radio programs, speeches, lectures, performances, language recordings
Video/Film	full feature films, documentaries, news clips, anthropological/expedition footage, home movies, animation
Graphics	3D models, simulated walk-throughs of buildings, archeological sites, VRML
Multimedia	presentations, slide shows, SMIL files, QuickTime VR

Table 1: Overview of Museum Multimedia Types

Often multimedia content has been generated for preservation and dissemination purposes. The multimedia resource may be a digital surrogate of the original culturally significant museum artifact which is too valuable or fragile to be handled or is inaccessible for reasons of location. In many situations, the multimedia resources are created as an alternative visual representation (image, model) of a physical museum artifact or as a replacement for earlier analog and digital formats which are becoming obsolescent.

Hence in addition to the typical bibliographic information, the metadata for multimedia resources may need to describe detailed formatting information, structural or segmentation information (temporal, spatial and spatio-temporal segments), semantic information (description of the objects/people/places/events which are recorded) and the event history and rights information. A detailed description of the metadata requirements for multimedia in museums is provided in the next section.

2.2 The Metadata Requirements for Multimedia in Museums

The metadata associated with multimedia resources can be classified into a number of different categories:

- **Bibliographic metadata** - this includes information about the resource's creation/production (date, place) and the individuals or organizations involved (e.g., producer, director, and cast) and the resource's classification information (e.g., title, abstract, subject, and genre).
- **Formatting metadata** - this includes information about the format, encoding, storage and system requirements associated with the resource.

80

image	audio	video	text	multimedia
format (image/tiff)	format (audio/aiff)	format (Quicktime, MPEG1)	format (text/ms word)	format (text/html)
filesize (bytes)	filesize (bytes)	filesize (bytes)	filesize (bytes)	filesize (bytes)
version (v 4.0)	version (v2.5)	version (v1.1)	version = 97	version (v3.0)
resolution (600dpi)	samplingrate (44.1kHz)	dimensions (640x480)	compression (zip)	software (MS FrontPage 2.0)
dimensions (1024x768)	samplesize (16 bit)	aspectratio (4:3)	characterset (Unicode)	storagetype (HD server)
aspectratio (4:3)	duration (04:45:56.34)	duration (32min 12sec)	template (summary.dtd)	template (Program.xsl)
colourdepth (8-bit grayscale, 24-bit colour)	compression (MPEG2/Layer 3)	compression (mp2)		bandwidth requirements
colourpalette (CMYK, RGB, GrayScale)	encapsulation (RealAudio G2)	encoding (mp2)		system requirements (OS, software, hardware,peripherals)
framerate (25fps)	tracks (mono, stereo)	sound (Yes/No)		
colourLUT (base64)	storagetype (Phillips DAT)	storagetype (DVD)		
orientation (Portrait, Landscape)		colour (Colour or B/W)		
compression (CCIT4)		special Effects (ChromaKey)		
storagetype (CD-ROM, Jazz, hard drive)		delivery and presentation requirements (bandwidth, operating system, hardware,)		
scanner (make, model, serial #)		software, peripherals		
		camera details and settings (make, model, serial#, aperture, focallength, filter)		

Table 2: Typical Format Metadata for Different Media Types

Table 2 shows the formatting metadata typically recorded for different media types.

- **Structural metadata** - this provides information about the structural decomposition of the multimedia resource into spatial, temporal or spatio-temporal segments (scenes, shots, frames, image regions) and the relationships between these segments.
- **Content metadata** - this provides indexes to the actual content which is recorded or depicted within the multimedia resource. Content metadata can vary from natural language descriptions of the people, objects, places or events which are depicted to the low level audio or visual features such as colour histograms or volume.
- **Events and rights metadata** - this is information describing the life history of the resource. It includes everything from acquisition and relocation events to the reformatting, editing, repackaging and distribution events to the metadata attribution events to the usage, copyright agreements, and permission events.

A number of projects have developed or are developing metadata models for multimedia in museums (Gabriel, 2001; The TOKEN 2000) or for historical audiovisual collections (The ECHO). These projects are either developing their own application-specific data models and vocabularies, or choosing one of the existing standards (MPEG-7 or CIDOC/CRM). None have considered the approach of merging ontologies from the museum domain and the multimedia domain into a single ontology.

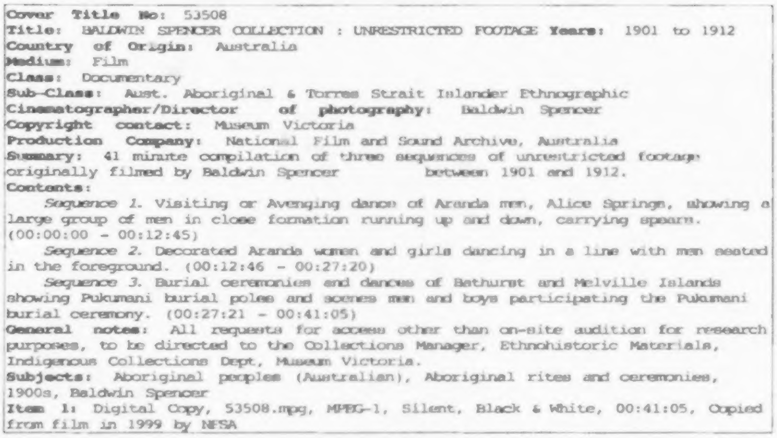

Cover Title No: 53508
Title: BALDWIN SPENCER COLLECTION : UNRESTRICTED FOOTAGE **Years:** 1901 to 1912
Country of Origin: Australia
Medium: Film
Class: Documentary
Sub-Class: Aust. Aboriginal & Torres Strait Islander Ethnographic
Cinematographer/Director of photography: Baldwin Spencer
Copyright contact: Museum Victoria
Production Company: National Film and Sound Archive, Australia
Summary: 41 minute compilation of three sequences of unrestricted footage originally filmed by Baldwin Spencer between 1901 and 1912.
Contents:
Sequence 1. Visiting or Avenging dance of Aranda men, Alice Springs, showing a large group of men in close formation running up and down, carrying spears. (00:00:00 - 00:12:45)
Sequence 2. Decorated Aranda women and girls dancing in a line with men seated in the foreground. (00:12:46 - 00:27:20)
Sequence 3. Burial ceremonies and dances of Bathurst and Melville Islands showing Pukumani burial poles and scenes men and boys participating the Pukumani burial ceremony. (00:27:21 - 00:41:05)
General notes: All requests for access other than on-site audition for research purposes, to be directed to the Collections Manager, Ethnohistoric Materials, Indigenous Collections Dept, Museum Victoria.
Subjects: Aboriginal peoples (Australian), Aboriginal rites and ceremonies, 1900s, Baldwin Spencer
Item 1: Digital Copy, 53508.mpg, MPEG-1, Silent, Black & White, 00:41:05, Copied from film in 1999 by NFSA

Example 1

2.3 A Typical Example

Consider an example typical of multimedia content held by museums or archives: a film owned by the Museum of Victoria which contains unedited footage of Australian Aboriginal tribal ceremonies filmed by anthropologist Baldwin Spencer between 1901 and 1912. Below is the catalogue item from ScreenSound Australia's online catalogue. In 1999, the original film was copied to digital format (MPEG-1) by the National Film and Sound Archive for the purpose of preservation.

In Sections 3.2 and 4.2 below we compare the abilities of the CIDOC CRM and MPEG-7 to describe museum multimedia, by describing this example using their domain-specific vocabularies.

3. The CIDOC/CRM

3.1 Overview of the CIDOC CRM

The "CIDOC object-oriented Conceptual Reference Model" (CRM), was developed by the ICOM/CIDOC Documentation Standards Group to provide an 'ontology' for cultural heritage information. Its primary role is to serve as a basis for mediation of cultural heritage information and thereby provide the semantic 'glue' needed to enable wide area information exchange and the integration of heterogenous resources between cultural institutions.

The CIDOC CRM is presented as an object-oriented extensible data model, expressed in RDF Schema (RDF Schema Spec 2000). Figure 1 illustrates the class hierarchy for the CIDOC CRM, as generated by the SIS knowledge base (ICS). The detailed specification of the CIDOC CRM Version 3.2 which includes detailed descriptions of the class and property definitions, hierarchies and relationships, is available from (Crofts, 2002).

In order to understand and evaluate the CIDOC CRM's ability to describe multimedia resources, we apply it to the example in Section 2.3.

3.2 A CIDOC CRM Description of the Example

Together with an analysis of the class and property hierarchies provided by the CIDOC CRM, this exercise reveals that the CIDOC CRM's strengths lie in its ability to describe:

- Identification information;
- Acquisition and ownership information;

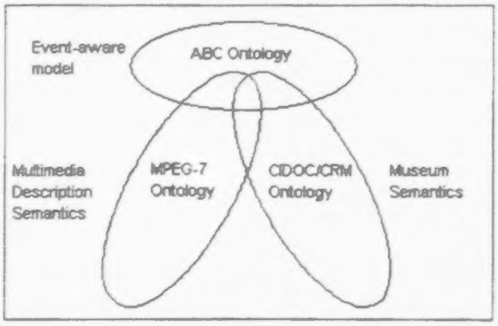

Figure 1 - The CIDOC/CRM Class Hierarchy

BEST COPY AVAILABLE

- Physical movement, location and relocation information;
- Physical attributes and features - dimensions, marks, visual items, material, sections, physical location etc.;
- Historical events - CIDOC CRM supports rich semantic descriptions of concepts or events - both real world events, as well as the events which occur in the life cycle of a resource, and those events which are depicted in the visual information objects.

However the CIDOC CRM is limited in its ability to describe digital objects and particularly digital multimedia or audiovisual content. The following requirements are inadequately supported:

- Formatting attributes (encoding, storage, system requirements) for digital images, audio, video, text and multimedia such as those shown in Table 1 are not currently supported;

- It is possible to define sections of physical objects using spatial measurements or coordinates and the temporal location of Events or Periods using the Date and TimeSpan entities - but not temporal, spatial or spatio-temporal locations within non-physical digital media;

- Physical features can be described but not visual or audio features such as colour histograms, regions, shape, texture, volume etc.;

- Hierarchical or sequential summaries of audiovisual content which specify keyframes, scene changes or key videoclips, are not supported

The CIDOC CRM provides the *is_documented_in* property to record the relationship of a culturally significant physical or real-world artefact or event to its visual/audio/audiovisual recording, which is classed as a *Document*. However there is no explicit support for the different media types in CIDOC/CRM. Because the CIDOC CRM is designed to provide an extensible underlying framework, it may be possible to improve support for multimedia descriptions, through the addition of MPEG-7 multimedia-specific sub-classes and sub-properties to existing CIDOC/CRM superclasses and super-properties. We investigate this approach in Section 5.

```
<?xml version="1.0" encoding="ISO-8859-1"?>
 <?xml-stylesheet type="text/xsl" href="crm.xsl"?>
 <CRMset>
  <CRM_Entity>
   <in_class>E22: Man-Made Object</in_class>
   <is_identified_by>Object ID 53508
    <in_class>E42: Object Identifier</in_class>
   </is_identified_by>
   <has_title>BALDWIN SPENCER COLLECTION : UNRESTRICTED FOOTAGE
     <in_class>E35: Title</in_class>
   </has_title>
   <has_type> Documentary Film
     <in_class>E55: Type</in_class>
   </has_type>
   <was_created_by> Baldwin Spencer
          <in_class>E21: Person</in_class>
     <took_place_at> Central Australia,
          <in_class>E53: Place</in_class>
          <falls_within> Australia
             <in_class>E53: Place</in_class>
          </falls_within>
     </took_place_at>
     <at_most_within>1901-1912
      <in_class>E52: Time Span</in_class>
     </at_most_within>
   </was_created_by>
   <has_current_owner> Museum of Victoria
     <has_type> Museum
     </has_type>
   </has_current_owner>
   <was_produced_by> National Film and Sound Archive
     <in_class>E21: Production Company</in_class>
   </was_produced_by>
   <used_general_technique>film
     <in_class>E55: Type</in_class>
   </used_general_technique>
   <depicts_concept>Aboriginal rites and ceremonies
     <in_class>E55: Type</in_class>
   </depicts_concept>
   <has_note>
    41 minute compilation of three sequences of unrestricted footage originally
    filmed by Baldwin Spencer between 1901 and 1912
     <has_type> Description
       <in_class>E55: Type</in_class>
     </has_type>
   </has_note>
   <has_note>All requests for access other than on-site audition for research
          purposes, to be directed to the Collections Manager, Ethnohistoric
          Materials,Indigenous Collections Dept, Museum Victoria.
   </has_note>
   <has_dimension>
    <name>duration</name>
    <value>00:41:05</value>
    <unit>minutes</unit>
   </has_dimension>
   <is_composed_of>
     Segment 1
    <has_type>videoSegment</has_type>
    <has_time_span>
      <begins_at>00:00:00</begins_at>
      <ends_at>00:12:45</ends_at>
    <has_time_span>
    <depicts_concept>Visiting or Avenging dance of Aranda men, Alice Springs,
        showing a large group of men in close formation running up and down,
        carrying spears
    </depicts_concept>
   </is_composed_of>
   <is_composed_of>
     Segment 2
    <has_type>videoSegment</has_type>
    <has_time_span>
      <begins_at>00:12:46</begins_at>
      <ends_at>00:27:20</ends_at>
    <has_time_span>
    <depicts_concept>Decorated Aranda women and girls dancing in a line with
         men seated in the foreground.
    </depicts_concept>
   </is_composed_of>
   <is_composed_of>
     Segment 3
    <has_type>videoSegment</has_type>
    <has_time_span>
      <begins_at>00:27:21</begins_at>
      <ends_at>00:41:05</ends_at>
    <has_time_span>
    <depicts_concept>Burial ceremonies and dances of Bathurst and Melville
        showing Pukumani burial poles and scenes of men and boys participating
        Islands in the Pukumani burial ceremony.
    </depicts_concept>
   </is_composed_of>
   <is_documented_in>53508.mpg
    <in_class>E31: Document</in_class>
    <was_created_by>National Film and Sound Archive
      <took_place_at> Canberra
         <in_class>E53: Place</in_class>
         <falls_within> Australia
            <in_class>E53: Place</in_class>
         </falls_within>
      </took_place_at>
      <at_most_within>1999
         <in_class>E52: Time Span</in_class>
      </at_most_within>
    </was_created_by>
   </is_documented_in>
  </CRM_Entity>
 </CRMset>
```

Example 2

BEST COPY AVAILABLE

4. MPEG-7 - the Multimedia Content Description Interface

4.1 An Overview of MPEG-7

The Moving Pictures Expert Group (MPEG), a working group of ISO/IEC, is expected to shortly release the final standard for MPEG-7 (MPEG-7 Overview), the "Multimedia Content Description Interface", a standard for describing multimedia content. The goal of this standard is to provide a rich set of standardized tools to enable both humans and machines to generate and understand audiovisual descriptions which can be used to enable fast efficient retrieval from digital archives (pull applications) as well as filtering of streamed audiovisual broadcasts on the Internet (push applications). MPEG-7 can describe audiovisual information regardless of storage, coding, display, transmission, medium, or technology. It addresses a wide variety of media types including still pictures, graphics, 3D models, audio, speech, video, and combinations of these (e.g., multimedia presentations).

MPEG-7 provides:

- a core set of Descriptors (Ds) that can be used to describe the various features of multimedia content;

- pre-defined structures of Descriptors and their relationships, called Description Schemes (DSs).

Initially MPEG-7 definitions (description schemes and descriptors) were expressed solely in XML Schema [9-11]. XML Schema proved ideal for expressing the syntax, structural, cardinality and datatyping constraints required by MPEG-7. However semantic interoperability is necessary to enable systems to exchange data (e.g., metadata descriptions), to understand the precise meaning of that data and to translate or integrate data across systems or from different metadata vocabularies. Hence it was recognized that there was a need to formally define the semantics of MPEG-7 terms; and to express these definitions in a machine understandable, interoperable language. RDF Schema (RDF Schema Spec 2000) was the obvious choice due to its ability to express semantics and semantic relationships through class and property hierarchies and its endorsement by the W3C's Semantic Web Activity (W3C). Consequently the Adhoc Group for MPEG-7 Semantic Interoperability was established and an MPEG-7 ontology was developed and expressed in RDF Schema and DAML+OIL extensions (Hunter Adding, Hunter An ROF 2001). The extensions provided by DAML+OIL (DAML+OIL) were necessary to satisfy certain requirements such as the need for multiple ranges and sub-class specific constraints. The basic class hierarchy of MPEG-7 content and segments is shown in Figure 2 e.g., the MPEG-7 class *VideoSegment* is a subclass of both *Video* and *Segment*.

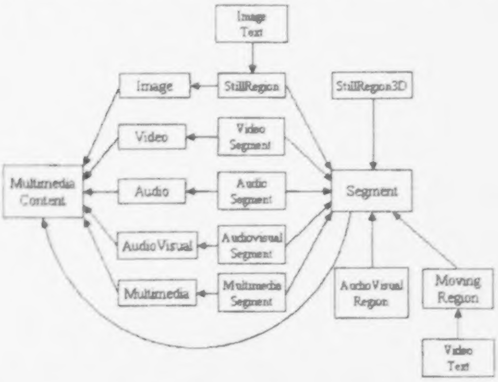

Figure 2 - MPEG-7 Multimedia Segment Class Hierarchy

The relationships of the Segment types to other segment types and multimedia entities are dependent on the allowed types of decomposition. Multimedia resources can be segmented or decomposed into subsegments through 4 types of decomposition:

- **Spatial Decomposition** - e.g., spatial regions within an image;
- **Temporal Decomposition** - e.g., temporal video segments within a video;
- **Spatiotemporal Decomposition** - e.g., moving regions within a video; or by
- **MediaSource Decomposition** - e.g., the different tracks within an audio file or the different media objects within a SMIL presentation

Associated with each of the subclasses in Figure 2 are various properties which define permitted relationships between the segment classes corresponding to specific structural or organizational description schemes and the permitted audio, visual and audiovisual attributes associated with different types of multimedia segments.

BEST COPY AVAILABLE

Type	Feature	Descriptors
Visual	Color	DominantColor
		ScalableColor
		ColorLayout
		ColorStructure
		GoFGoPColor (extension of ColorStructure)
	Texture	HomogeneousTexture
		TextureBrowsing
		EdgeHistogram
	Shape	RegionShape
		ContourShape
		Shape3D
	Motion	CameraMotion
		MotionTrajectory
		ParametricMotion
		MotionActivity
Audio	Silence	Silence
	Timbre	InstrumentTimbre
		HarmonicInstrumentTimbre
		PercussiveInstrumentTimbre
	Speech	Phoneme
		Articulation
		Language
	Musical Structure	MelodicContour
		Rhythm
	SoundEffects	Reverberation, Pitch, Contour, Noise

Table 3 - MPEG-7 Visual and Audio Features and their Corresponding Descriptors

The visual and audio features which may be associated with multimedia and segment classes are listed in Table 3. Associated with each of the visual and audio features is a choice of descriptors, also illustrated in Table 3. Precise details of the structure and semantics of these descriptors are provided in ISO/IEC 15938-3 FCD Multimedia Content Description Interface - Part 3 Visual and ISO/IEC 15938-3 FCD Multimedia Content Description Interface - Part 4 Audio (ISO, 2001).

Only particular visual and audio descriptors are applicable to each segment type. Table 4 illustrates the association of visual and audio descriptors to different segment types. The MPEG-7 RDF Schema (Hunter,Adding; Hunter,An ROF) specifies the constraints on these property-to-entity relationships.

Feature	Video Segment	Still Region	Moving Region	Audio Segment
Time	X	-	X	X
Shape	-	X	X	-
Color	X	X	X	-
Texture	-	X	-	-
Motion	X	-	X	-
Audio	X	-	-	X

Table 4 - Relationships between Segment Types and Audio and Visual Descriptors

In addition to the basic Multimedia and Segment entities and the visual and audio descriptors, MPEG-7 provides standardized Description Schemes which combine the base classes and properties above, into pre-defined structured relationships (ISO...-5, 2001. Figure 3 provides an overview of the organization of MPEG-7 Multimedia DSs into the following areas: Basic Elements, Content Description, Content Management, Content Organization, Navigation and Access, and User Interaction.

These different MPEG-7 DSs enable descriptions of multimedia content which cover:

- Information describing the creation and production processes of the content (director, title, short feature movie);
- Information related to the usage of the content (copyright pointers, usage history, broadcast schedule);
- Media information of the storage features of the content (storage format, encoding);
- Structural information on spatial, temporal or spatio-temporal components of the content (scene cuts, segmentation in regions, region motion tracking);

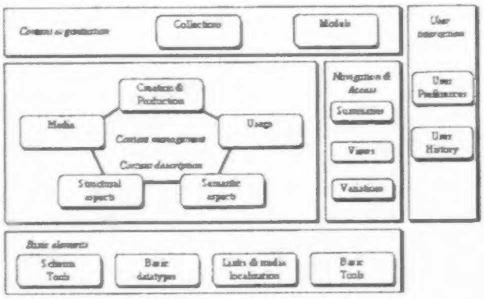
Content organization
Collections
Models
User interaction
Creation & Production
Media
Usage
Content management
Content description
Structural aspects
Semantic aspects
Navigation & Access
Summaries
Views
Variations
User Preferences
User History
Basic elements
Schema Tools
Basic datatypes
Links & media localization
Basic Tools

Figure 3 - Overview of MPEG-7 Multimedia DSs (ISO/IEC)

BEST COPY AVAILABLE

- Information about low level features in the content (colors, textures, sound timbres, melody description);
- Conceptual, semantic information of the reality captured by the content (objects and events, interactions among objects);
- Information about how to browse the content in an efficient way (summaries, views, variations, spatial and frequency sub-bands);
- Organization information about collections of objects and models which allow multimedia content to be characterized on the basis of probabilities, statistics and examples;
- Information about the interaction of the user with the content (user preferences, usage history).

In the next section we generate an MPEG-7 description of the example in Section 2.3, which illustrates the use of the CreationInformationDS, the MediaInformationDS and the TemporalDecompositionDS, to provide meaningful structure to the metadata description.

4.2 An MPEG-7 Description of the Example

Example 3 is an MPEG-7 description of the example in Section 2.3.

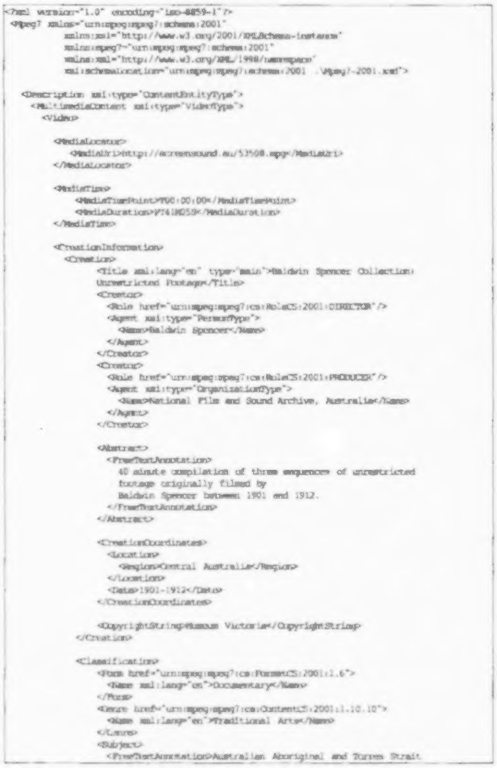

```
<?xml version="1.0" encoding="iso-8859-1"?>
 <Mpeg7 xmlns="urn:mpeg:mpeg7:schema:2001"
          xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
          xmlns:mpeg7="urn:mpeg:mpeg7:schema:2001"
          xmlns:xml="http://www.w3.org/XML/1998/namespace"
          xsi:schemaLocation="urn:mpeg:mpeg7:schema:2001 .\Mpeg7-2001.xsd">

  <Description xsi:type="ContentEntityType">
   <MultimediaContent xsi:type="VideoType">
    <Video>

     <MediaLocator>
       <MediaUri>http://screensound.au/53508.mpg</MediaUri>
     </MediaLocator>

     <MediaTime>
        <MediaTimePoint>T00:00:00</MediaTimePoint>
        <MediaDuration>PT41M05S</MediaDuration>
     </MediaTime>

     <CreationInformation>
      <Creation>
           <Title xml:lang="en" type="main">Baldwin Spencer Collection:
           Unrestricted Footage</Title>
           <Creator>
            <Role href="urn:mpeg:mpeg7:cs:RoleCS:2001:DIRECTOR"/>
            <Agent xsi:type="PersonType">
             <Name>Baldwin Spencer</Name>
            </Agent>
           </Creator>
           <Creator>
            <Role href="urn:mpeg:mpeg7:cs:RoleCS:2001:PRODUCER"/>
            <Agent xsi:type="OrganizationType">
             <Name>National Film and Sound Archive, Australia</Name>
            </Agent>
           </Creator>

           <Abstract>
            <FreeTextAnnotation>
             40 minute compilation of three sequences of unrestricted
             footage originally filmed by
             Baldwin Spencer between 1901 and 1912.
            </FreeTextAnnotation>
           </Abstract>

           <CreationCoordinates>
            <Location>
             <Region>Central Australia</Region>
            </Location>
            <Date>1901-1912</Date>
           </CreationCoordinates>

           <CopyrightString>Museum Victoria</CopyrightString>
      </Creation>

      <Classification>
           <Form href="urn:mpeg:mpeg7:cs:FormatCS:2001:1.6">
            <Name xml:lang="en">Documentary</Name>
           </Form>
           <Genre href="urn:mpeg:mpeg7:cs:ContentCS:2001:1.10.10">
            <Name xml:lang="en">Traditional Arts</Name>
           </Genre>
           <Subject>
            <FreeTextAnnotation>Australian Aboriginal and Torres Strait
```

Example 3

5. Comparison of the CIDOC/CRM and MPEG-7 Ontologies

5.1 Overlaps, Intersections and Differences

A comparison of the MPEG-7 and CIDOC CRM ontologies and their descriptions of the same resource above, reveals the following:

- Both metadata models are capable of describing the creation, production and classification information associated with a resource. Mappings between these components of the two models is possible;

- The CIDOC CRM is more focussed on describing physical museum artefacts and real world events from an epistemiological perspective. CIDOC CRM provides an ontology which allows the decomposition of knowledge available in data records into atomic propositions that are context free, interpretable when stand-alone but can easily be compiled into an integrated knowledge base;

- MPEG-7 is more focused on precise, fine-grained content-based descriptions of multimedia content, particularly digital multimedia, to enable the automated search and retrieval or filtering and retrieval of relevant multimedia content using standardized descriptions;

- Because the CIDOC CRM vocabulary is based on a hierarchical object-oriented model in which *Events* and *Activities* are core entities, it provides a better underlying framework for recording the events, changing attributes and dynamic relationships associated with a resource.

BEST COPY AVAILABLE

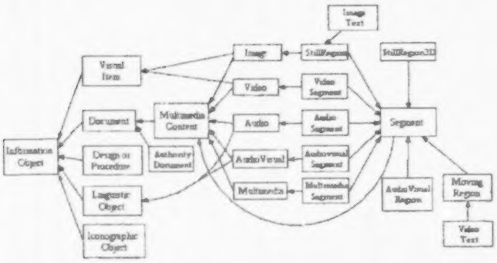
Image Text
Visual Item
Image
StillRegion
StillRegion3D
Video
Video Segment
Document
Multimedia Content
Audio
Audio Segment
Segment
Information Object
Design or Procedure
Authority Document
AudioVisual
Audiovisual Segment
Linguistic Object
Multimedia
Multimedia Segment
AudioVisual Region
Moving Region
Iconographic Object
Video Text

Figure 4 - Extending the CIDOC CRM Class Hierarchy with MPEG-7 Multimedia Subclasses

5.2 Merging the Ontologies

Since CIDOC CRM is designed to provide a top level set of classes and properties which can act as attachment points for domain-specific metadata ontologies, it makes sense to use the CIDOC CRM as the foundation and to extend it with MPEG-7 specific components, to add multimedia metadata capabilities. The obvious attachment point for the MPEG-7 class hierarchy is the CIDOC CRM *Document* class, as shown in Figure 4.

The CIDOC CRM provides an *is_composed_of* property which can be extended through RDF Schema sub-properties to define the structural or segmentation metadata associated with multimedia resources. Spatial, temporal, spatio-temporal and media-source decompositions are all provided through this approach, as illustrated in Figure 5.

The CIDOC CRM *Time-Span* and *Place* classes need to be sub-classed to enable MediaTime and MediaPlace classes to be defined - these are required in order to specify temporal, spatial and spatio-temporal locators within audiovisual resources.

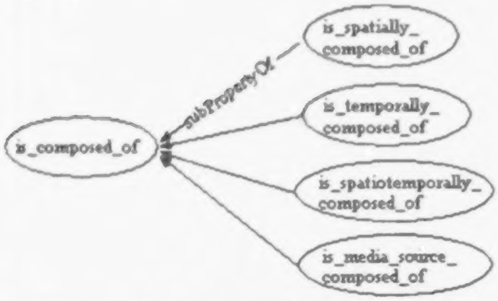
is_spatially_composed_of
subPropertyOf
is_temporally_composed_of
is_composed_of
is_spatiotemporally_composed_of
is_media_source_composed_of

Figure 5 - Adding MPEG-7 Segmentation properties to CIDOC-CRM

```
<has_format>
      <name>MimeType</name>
      <value>video/mpeg-1</value>
   </has_format>

   <has_format>
      <name>frame_rate</name>
      <value>25</value>
      <unit>fps</unit>
   </has_format>
```

Example 4

Further extensions to the CIDOC CRM which are required for adequate multimedia description include the provision of formatting properties (Table 2) and visual and audio features/descriptors (Table 3), associated with multimedia and segment entities.

Some formatting metadata can be accommodated by the existing *Dimension* class (e.g., file size, frame dimensions). But formatting information such as encoding and storage medium could be attached through a new *Format* class and a new *has_format* property, similar to the existing *Dimension* class as shown in Example 4.

Audio and visual descriptors could be provided through the provision of two new classes, *VisualFeature* and *AudioFeature*. Their respective MPEG-7 descriptors as outlined in Table 3 could be defined as sub-classes. The properties *has_visual_feature* and *has_audio_feature* would be required to associate these new classes/sub-classes to the relevant Multimedia Document types, as specified in Table 4.

Based on the proposals described here, an RDF Schema representation of the MPEG-7 extensions to the CIDOC CRM, has been developed and is available (RDF Schema Rep).

Feature	Video Segment	Still Region	Moving Region	Audio Segment
Time	X	-	X	X
Shape	-	X	X	-
Color	X	X	X	-
Texture	-	X	-	-
Motion	X	-	X	-
Audio	X	-	-	X

Table 4

BEST COPY AVAILABLE

6. Conclusions and Future Work

6.1 Conclusions

In this paper we have analyzed the strengths and weaknesses of the CIDOC CRM and MPEG-7 ontologies in the context of providing a metadata model for describing and managing museum multimedia resources. Based on this analysis, we have described an approach to merging the two models by using the CIDOC CRM as the underlying foundation and extending it through the addition of MPEG-7-specific sub-classes and sub-properties to provide support for multimedia concepts and descriptions. The outcome is a single machine-understandable ontology (in RDF Schema) which can be used to provide the underlying model for describing multimedia in museums and thus facilitate the exchange, sharing and integration of heterogeneous multimedia information between cultural institutions.

6.2 Future Work

Having developed the model, the next step is to test, evaluate and refine it by applying it to the description and management of real collections of multimedia resources within museums. We plan to do this through the development of a test-bed using resources provided by the Smithsonian's National Museum of the American Indian CRC.

In 2001, two workshops were held by the DELOS Working Group on Ontology Harmonization (Doerr) to discuss and compare the CIDOC CRM and ABC ontologies (Lagaozi, Hunter, 2001). Both of these ontologies have been designed to facilitate semantic interoperability between metadata vocabularies from different domains. A third workshop is planned for June 2002, and the anticipated outcome will be a common merged model. Assuming that this eventuates, then future work will involve determining how the multimedia specific concepts of MPEG-7 can be attached to or accommodated within this merged ontology.

Acknowledgements

The work described in this paper has been carried out as part of a Smithsonian Fellowship funded by the Queensland Government and the Smithsonian Institute. Additional funding has been provided by the Cooperative Research Centre for Enterprise Distributed Systems Technology (DSTC) through the Australian Federal Government's CRC Programme (Department of Industry, Science and Resources). Thanks also to the other members of the DELOS Working Group on Ontology Harmonization, and especially Martin Doerr, both for Figure 1, and for his patient and detailed explanations of the CIDOC CRM.

References

CIDOC Conceptual Reference Model, http://cidoc.ics.forth.gr/

Crofts, N. Dionissiadou, I., Doerr, M., Stiff, M. (editors), "Definition of the CIDOC object-oriented Conceptual Reference Model , Feb 2002 (version 3.2.1) http://cidoc.ics.forth.gr/docs/cidoc_crm_version_3.2.1.rtf

DAML+OIL, March 2001. http://www.daml.org/2001/03/daml+oil-index

Doerr, M. DELOS Working Group on Ontology Harmonization, Workshop Reports, http://www.metadata net/harmony/Collaboration.htm

The ECHO European CHronicles On-line Project http://pc-erato2.iei.pi.cnr.it/echo/

Gabriel, D., Ribeiro, C. "A Metadata Model for Multimedia Databases", ICHIM 01, Milan, September 2001 http://www.archimuse.com/ichim2001/abstracts/prg_100000544.html

Hunter, J. "Adding Multimedia to the Semantic Web - Building an MPEG-7 Ontology", http://archive.dstc.edu.au/RDU/staff/jane-hunter/swws.pdf

Hunter, J. "An RDF Schema/DAML+OIL Representation of MPEG-7 Semantics", MPEG, Document: ISO/IEC JTC1/SC29/WG11 W7807, December 2001, Pattaya

ICS FORTH, "The Semantic Index System (SIS)", http://www.ics.forth.gr/proj/isst/Publications/paperlink/SIS-descr.ps.gz

ISO/IEC 15938-5 FCD Information Technology - Multimedia Content Description Interface - Part 5: Multimedia Description Schemes, March 2001, Singapore

ISO/IEC 15938-3 FCD Information Technology - Multimedia Content Description Interface - Part 3: Visual, March 2001, Singapore

ISO/IEC 15938-4 FCD Information Technology - Multimedia Content Description Interface - Part 4: Audio, March 2001, Singapore

Lagoze, C., and Hunter, J. "The ABC Ontology and Model", Journal of Digital Information, Volume 2, Issue 2, November 2001, http://jodi.ecs.soton.ac.uk/Articles/v02/i02/Lagoze/

MPEG-7 Overview, http://mpeg.telecomitalialab.com/standards/mpeg-7/mpeg-7.htm

RDF Schema Representation of the MPEG-7 Extensions to the CIDOC CRM, http://metadata.net/harmony/MPEG-7_CIDOC_CRM.rdfs

RDF Schema Specification 1.0, W3C Candidate Recommendation 27 March 2000. http://www.w3.org/TR/rdf-schema/

The TOKEN2000 project http://dbs.cwi.nl:8080/cwwwi/owa/cwwwi.print_projects?ID=78

W3C Semantic Web Activity http://www.w3.org/2001/sw/

XML Schema Part 0: Primer, W3C Recommendation, 2 May, 2001, http://www.w3.org/TR/xmlschema-0/

XML Schema Part 1: Structures, W3C Recommendation, 2 May 2001, http://www.w3.org/TR/xmlschema-1/

XML Schema Part 2: Datatypes, W3C Recommendation, 2 May 2001, http://www.w3.org/TR/xmlschema-2/

Today's Authoring Tools for Tomorrow's Semantic Web

Andy Dingley, Codesmiths, Bristol and Paul Shabajee,
Graduate School of Education & Institute for Learning and
Research Technology, University of Bristol,UK

Abstract

This paper reports on the development of a prototype authoring tool developed as part of on-going research around the needs of the ARKive project. The project holds text, rich-media and descriptions of factual statements about biodiversity and conservation information. A key user community is that of school age children, requiring the mark-up of educational metadata in open standards such as IEEE LOM. A previous paper by the authors reported on the publishing architecture for this project. This publishing architecture is intended to serve a range of audiences (ages, language and level of language skills). By storage of the content as discrete units, with extensive metadata describing each one, units may be retrieved and served to the audience as appropriate. Future developments may extend this to support ad hoc queries, not just rigidly pre-defined standard pages.

Authoring development has shown that a simple and pragmatic tool based on Microsoft Word may still address advanced technologies such as RDF, DAML and the future of the Semantic Web. Careful design has separated the process of describing a museum's exhibits, and the problem domain of the museum's area of interest. This gives two advantages. First, most of the effort now supports a generic on-line museum that may be re-targeted from bio-diversity to any other topic. Secondly, solving the problem domain by ontological descriptions, not rigid program code, gives the ability to easily reference pre-existing or external vocabularies. This improves the flexibility of solving the initial problem, allows the same code to be re-used on other projects, and assists publishing into other metadata formats

Keywords: Semantic Web; RDF; PAML; Metadata; Dublin Core; Knowledge Representation; Biodiversity Information; ARKive

The ARKive Project

ARKive (www.arkive.org) is a Wildscreen Trust (www.wildscreen.org) initiative to build a Web-based multimedia digital archive of the world's endangered animals and plant species. The project brings together thousands of films, videos, sounds and photographs of threatened and recently extinct species. Hewlett Packard Laboratories (http://www.hpl.hp.com/arkive/) are supporting ARKive by funding a research team to develop the technical infrastructure, including the content management software.

The effective and efficient educational use of these materials is fundamental to the aims of the project. A previous paper (Dingley and Shabajee, 2001) described a publishing architecture to offer content tailored to the dynamic needs of multiple educational user groups. This paper describes the on-going development of a prototype-authoring tool to provide content to support this.

This paper reports on the development of a prototype-authoring tool developed as part of on-going research around the needs of the ARKive type projects by the authors. The work reported here is taking place independently of, and in parallel with, the development of the ARKive project by the Wildscreen Trust.

Project Teams

ARKive and similar projects incorporate many groups, each with its own agenda for the content authoring problems. The film-makers have little interest in authoring, but need shot-logging of the existing film library items (the process of describing scenes and shots on a piece of film and the objects they depict, often to extreme detail). Educators are interested in editing multiple copies of the same content, re-written for different audiences (language, content target age, language skills). Scientists

BEST COPY AVAILABLE

require statements of complex facts in a way that is machine-accessible to allow comparison between species, and that may be published to other sites. The Semantic Web (Berners-Lee, 1998) researchers were interested in the techniques, more than the content.

The Publishing Architecture

The core of the prototype publishing architecture is that of dynamically assembling the publication from a large set of content units, each identified by detailed metadata. Some content units contain knowledge of different topics; others are the same material re-formatted for a different audience.

Content units

A content unit may contain different media types: plain text, machine-processable expressions of facts, images, video, as examples, although this list is deliberately open-ended for any new formats. All content units appear identical to the data store; they are stores of opaque content, with an associated transparent set of metadata. This metadata describes the application content ("this is a lion", "feeding behavior is shown"), the technical format ("this is a QuickTime video"), and a great deal of additional description ("the narration is in English", "the scene of the wildebeest being eaten is not suitable for 7-year -olds", "the copyright owner is…"). For many text units, the metadata is several times the size of the content itself – this is by deliberate design, and in any case, all space concerns are dwarfed by the video.

The scope of a content unit may be variable. In the simplest case, for ARKive, each unit contains a "species page"; a complete description of one species for one audience. This is equivalent to the current Web site; a traditional database-backed collection of static pages. The next stage is to provide content units for each section of the page (appearance, habitat, distribution etc.) and to duplicate these units for each audience and language combination. Producing a page is now a filtering operation; the relevant set of units is retrieved for that species and then filtered to choose the most appropriate unit of each set for the target audience. The set of units may itself be filtered: a general interest audience might not receive some detailed scientific content, a limited browser device might have images but not video, and young children might be spared some carnivorous images. Overall ratings e.g. PICS (http://www.w3.org/PICS/) or IEEE LOM (IEEE LTSC, 2001) may be aggregated from those stored on each content unit.

Authoring and Storage

Content units are authored as part of a species-specific document, as this is the way the researchers and authors work. The species documents produced by the authoring tool reflect this scope. Each contains a simple header of workflow information, and then a list of content units. There is no implied structure of these units; they do not need to follow the paragraphs or section headings in the anticipated published species page.

"Research notes" may be stored at the level of the species document, or at each content unit. These are very simple free-text notes, intended for the author's guidance, and are never published. Our authoring workflow must cope with multi-author authoring, typically where a species expert may write a terse scientific description which a content author then re-writes or extends for a target audience. Writing appropriately for younger audiences, or translation, are expected use cases of this.

Content units are to be stored in a large RDF data store. The development of such stores is an area of current research, and so for pragmatic reasons we have deferred its implementation. In the future, each new species loaded will become part (although still identifiable) of this greater whole. At present we keep each species document as a separate XML file.

The future of ARKive may move away from this dependency on the species pages. Content might also be published as a habitat description (authored for one species, and then made available to all that share that habitat), a description of rare or endangered geology, or a cross-species description of British woodland carnivores.

Factoids

Early work in Web knowledge bases expressed descriptions of knowledge in a human readable format. The next step was to publish this knowledge in a way that made it machine accessible. This may be termed the "bottom up" approach (Motta, Buckingham Shum & Domingue, 2000). The contrary "top down" approach first produces a sufficiently expressive data model to meet the consumer's needs, distributes this by manual or automatic means, and then describes the content so as to meet the data model. As interested parties can now either use these shared standards, or at least transform their own representations in and out of it, this can form the basis of a Semantic Web. Dublin Core (http://purl.org/dc) is an example of this approach. Although sometimes dismissed as a "mere" lowest common denominator, widespread Dublin Core would still represent a considerable advance over most current practice.

Real progress in more intelligent ARKive type sites also requires what we have termed "factoids", an internal expression of an external fact. Although common standards for exchanging data can improve interoperability, a site that wishes to use this information to drive searches must also have some level of understanding about its meaning – an ontology (see below).

In the case of ARKive, the data model represents its knowledge as a large set of content units, which may each contain factoids. Factoids may be seen as workflow units, allowing their source, completion and validation status to be tracked. They may also refer to a property whose meaning is defined in an ontology (see below). This meaning may only need to be opaque, but distinctly identifiable, i.e. it is still useful to recognize properties that are meaningful to compare, even if there is no machine understanding of their meaning. Our experience is that a factoid-based solution has these advantages over a text-based solution:

- Searching and general machine processing. There's a limit to what is possible with simple free-text searching, and any knowledge base as interesting as ARKive is far beyond this.

- Reification (making statements about statements), which offers the opportunity for validation and maintenance.

- Keeping track of where facts are referred to from other texts or media; e.g. species information or narration for a piece of video.

Content units contain factoids. A factoid is ARKive's expression of a fact, which is an externally pre-existing component of knowledge. A factoid will make reference to an ontology, and to be useful to the world outside our project, this ontology must be expressed in a communicable format, such as DAML+OIL (http://www.daml.org/).

Ontologies

An ontology is a formalized description of classes (things) and their related properties (statements which may be made about these things). The simplest level of processing merely identifies these properties, so that related properties may be automatically recognized as comparable. Apples and oranges both have a comparable property for "country of origin", but "number of segments" only applies to one. More sophisticated reasoning may allow inferencing that oranges and lemons are both citrus fruit, and so lemons may also have this property.

Use of Ontologies

The prototype uses DAML+OIL to represent its ontologies. They are used in two places. One represents the content unit and factoid structure. The other (of which there may be multiple instances) represents the external and publishable definitions of facts used by the factoids.

Content Store

The structural ontology for the content units is primarily a data-modeling exercise and would be familiar to any software developer with a background in object orientation or relational databases.

Descriptive Vocabularies

The descriptive vocabularies are simple in structure and mainly contain vocabulary lists. Vocabulary items may themselves have structure; an identifier unique within that vocabulary, a title (which may be repeated in other languages) and an optional description of their appropriate meaning.

92

Some of these vocabularies (target audience, target language, rich media format) are likely to be created and controlled by ARKive. They may also contain representations of each item in other well-known vocabularies; i.e. some of the ARKive audiences also contain their equivalents in IEEE LOM. By presenting the authors with a choice from a single short and familiar list of target audiences, kept within the control of the project, they avoid having to make decisions about target ages based on a foreign country's school system.

The authoring tool loads vocabularies dynamically. A file directory is searched at start-up, and all those found therein are loaded. Authors may then select any of these to describe each factoid.

The Flying Bat Problem

There is a third use for ontologies within ARKive type systems: inferencing.

Searching ARKive for all the flying vertebrates should return both birds and bats. In simple implementations, this is likely to fail. Bats will be returned, because their flying behavior is unusual for a mammal and so is explicitly stated. For birds, flying is too trivial to describe and so it will probably never have been expressed in a machine searchable form. Even if birds' flight were to be stated for all birds, this would represent an enormous expansion of data volume and authoring work.

A common solution to this is to list a long string of unstructured keywords, hoping that they include all such concepts. The problem then is that although this improves the situation slightly, by offering a low-cost means of stating the trivial facts, it is so vague, unmanageable and restricted as to be near-useless. An obvious example is that of penguins. Such a simple list can express set membership, but not set exclusion. Stating "doesn't fly" for penguins, then querying with a simple text search will now return the flying vertebrates as being bats and penguins, but will still exclude eagles.

The power of an ontology is that it allows automatic inferencing to solve this. It is possible to state simultaneously that "birds fly", "penguins don't fly" and "mammals don't fly", and for a suitable reasoner to then determine correctly that "bats and birds other than penguins can fly", based on this ontology and an authored statement describing bats, the exception. In a sufficiently large knowledge domain, this ability to generalize and infer is essential, if the authoring requirement is not to be impossible.

Inferencing may also require an audience-related qualification. A search for "large fish" should exclude whales and dolphins, being marine mammals, but should this still be true for a 6-year-old Captain Ahab who is unaware of the distinction? This topic of ultimate accuracy over understanding is one of ongoing debate with our educational experts, although the technical team is still keen to represent it, as a technology demonstrator.

Ontology Tools

Ontology authoring is complex, and assisted editing tools exist for it. We evaluated both Protégé (http://protege.standford.edu) and OilEd (http://www.ontoknowledg.org/oil). Sadly, the standards for our chosen ontology description language (DAML+OIL) are changing in advance of the tool support, and so we found that reverting to a simple text or XML editor was necessary to use all of its features.

Protégé and OilEd have been developed by groups from different communities. Protégé's roots are in object and database modeling. It is easy to use for ontologies (like our content store) that approximate this class of problem. OilEd arises from the knowledge representation and reasoning field. We found it harder to use for the simpler ontologies, but more appropriate for the inferencing problems.

Nominals

An early prototype of the content and fact store was built. To avoid digression when species experts quibbled over our scientific data, rather than our techniques for storing it, we avoided biological data for real species and chose to describe Pokémon instead. The complete dataset is also conveniently available. This turned out to be quite a different problem from describing our real-world species, much simpler, yet illuminating.

A "hard" problem in knowledge representation is that of nominals. These are sets of values that are somewhere between the ontology and the instance data. An ontology might have a formalized representation of "fish" and "fowls" as distinct classes, and of color as a property, but it does not describe a set of "blue birds" or "silver fish". In many instances though, it is necessary to reason about these sets just as if they were described as classes in the ontology.

For Pokémon, there were no nominals. All of the descriptive properties that were required, and would ever be required, were clear to us from the outset. The Pokémon ontology could thus describe all of the necessary classes itself. For our real-world data, we continually encountered emergent properties: ways of structuring the data that were not apparent until authoring its content itself. Some antelope have a style of movement called "pronking" (trotting with all four feet moving together). Describing one antelope that pronks is simple (text will suffice), but a second pronking antelope should use a categorization factoid for its movement that is identical (not just a matching word in free-text), even though this term does not appear in the vocabulary.

Interoperability

Interoperability with other sites turns out to be relatively simple, compared with the complexity of our own internal authoring process. Our embedded metadata publishing used Dublin Core and extensive use of its qualifier mechanism. Where suitable established standards existed (e.g. IEEE LOM) we stored their equivalent terms within our own vocabulary lists, then published the well-known term in preference to our internal identifier.

For the publication of metadata from a Web site, a small amount of knowledge and effort can soon bring a site to a level in advance of likely practice for the near future. Existing good practices such as publishing stable URL's, avoiding meaningless identifiers in URL's, etc., are still as valuable as they are for static sites.

Implementation

What we didn't do

Flat Text

Although ARKive intends to serve content in many final forms, most of them are heavily textual. The existing database-backed site simply holds a copy of the HTML code for each page, but in a SQL database. This is not using the power of the database as anything other than a content management system.

Issues of data accuracy were not a major drawback to flat text. Although simple typing errors are an obvious problem, the real problems are caused by subtle semantic errors, not simple syntax (e.g. issues related to interindexer consistancy as described in Markey 1984). This is particularly the case for judgment calls during video shot-logging: one person's notion of a "long shot" may be another's "close-up"

A major factor against the use of plain text was the development effort it would still require. It was always accepted that the major effort would need to be in supporting the authoring effort of selecting appropriate terms from large controlled vocabularies, particularly those for systematics and species taxonomy (Biosis 2002). These would always require authoring tools, no matter what the final format, and so the savings owing to a simplified format became proportionately less.

As may be seen later, the path we did choose fits in well with the integration of large amounts of preexisting flat text. The existing ARKive site already contains data and text on >100 species, and a legacy integration path had always been required. The main issue with importing legacy data turned out to be that of data quality, particularly in mapping informal taxonomic labeling onto a more rigidly structured vocabulary.

The DTD approach

In this scenario, a purely XML approach would have been taken. An initial DTD or XML Schema would have been produced, describing the data model used to represent the content.

Within this scenario, there are two possibilities: one based on a DTD and one on a Schema. Even though XML Schema is now a long-established standard, it is still ignored by the majority of DTD-based authoring tools. XML Schema was an attractive basis to build the future of the project around, remembering that this is still early days for a hopefully long-term archive, but obsolescent DTD's were certainly not.

94

Existing XML authoring tools were considered, but rejected for their dependencies on DTD's. This approach would probably have been the quickest from the viewpoint of development time and just-sufficient integration.

Custom Code

Writing our own authoring tools from scratch, probably in Java, would have been the most flexible. This could have used the many pre-existing XML tools, including RDF parsers such as ARP (Carroll, 2001). Java would, however, have required more effort to build an editor, slightly more complexity for deployment ,and, most importantly, it would have been an unfamiliar editor for the authors. There was also a scheduling problem, in that it would not have been available for authoring, however simple, until almost all of it was completed.

Platform

The chosen platform was that of Microsoft Word and its Visual Basic for Applications programming language, with the Microsoft Windows XML component. This offered a reasonably competent coding language, no need to code a text editor, and (most importantly) immediate familiarity for the content authors. None of Word's built-in HTML or XML features was used, as Word's view of appropriate XML use is not quite the same as that of other workers in the field. There are no pre-existing RDF tools accessible from this platform.

The code, and its integration with the editor, is very simple. Normally the author sees only a standard version of Word, with a template containing custom styles. To insert specific items (new blocks of content, links to rich media, machine-processable facts) a toolbar button is pressed, a dialog allows appropriate values to be browsed, and a block of Word-formatted text is then inserted. Word's styles are used to identify properties within the data model. To export the completed document, a macro then walks through the Word document, translating each paragraph and appending it to the XML DOM component, then finally saving the resultant XML.

An advantage of this approach is for project scheduling, always a problem for software development. It is very simple to generate simple documents, thus allowing early testing by users. Most of the development complexity and effort goes into adding the more sophisticated features, which may grow incrementally as they are coded.

The exported document is in RDF, serialized as simple XML. Experience with the video shot-logging sub-project showed that RDF is transformable by XSLT, although this is a painful process! An RDF data model representing the same content may be serialized to a number of valid XML representations, yet XML tools (Schema or XSL) have no conception that these are equivalent. Although workable XSL stylesheets may be coded, their reliability depends on their author having manually foreseen and coded for every valid variation in the representation of RDF. This complexity echoes the experience of other groups (Cawsey, 2000). In an early and optimistic phase of the project, it was thought that automatic tools could be developed to produce these tools in turn. This proved not to be the case.

Why XML is a given

One issue that did not arise as a point for discussion was that of using XML. It now seems that XML is the *sine qua non* of such applications, with no obvious competition. This was interesting, as the team also included at least one person with a background in SGML.

The main reason for choosing XML, particularly over SGML, is related to the availability of tools, rather than the qualities of the format itself. A consistent DOM is available for XML, from several environments. Although true code portability is still something of a pipe dream, this allows the Java developer community to at least hold a conversation with the Word developers.

We chose to represent the authored content with RDF, primarily because this is the technical team's core research interest, but it is also a useful format. RDF tools are still immature and offer little immediate advantage, but by careful design (primarily control of the serialization into XML) it is possible to treat RDF as if it were simply XML. If using RDF had any major costs associated with it, i.e. the simple XML / XSLT approach would have suffered, then we would have chosen not to use it. Creation of RDF is a simple matter, even with a purely XML toolset, but importing valid RDF from other sources is not a practical proposition without either a genuine RDF parser, or by imposing additional constraints on how the RDF is represented in XML.

Software Availability

The authoring software developed as part of this project has been open-sourced and made available (http://www.xcml.org/docbadger/).

Observations on External Standards

Vocabularies

Few of the vocabularies and thesauri we needed were available in any formal notation. Those that were are mainly the educational vocabularies, e.g. LOM.

Several of the subject-specific vocabularies, particularly for animal taxonomy, conservation status (http://www.wcmc.org) etc. were available in text form, with good provenance and stability. Various Web-scraping scripts were able to transform these into RDF or DAML documents. Other projects, such as Tim Berners-Lee's Semantic Web Road map (1998, http://www.w3.org/DesignIssues/Semantic.html) have taken a similar approach.

Property Sets

There is still no common adoption of property sets, for many common tasks. Dublin Core is an obvious solution to many of these, but there is scope for much improvement here. A surprising omission was bibliographic references, surely both a commonplace and easily formalized task. The de facto standard BibTex (http://www.isi.edu/webscripter/bibtex.o.daml) is still focused on solving the typographical problem of producing one's own papers, not the interoperability problem of sharing with others. While some groups wrestle with fundamentally difficult problems, a great advance towards a Semantic Web could be made just by wider adoption of the simplest steps.

Conclusions

- It is possible to develop useful authoring tools today, for full exploitation by tomorrow's Semantic Web. This is so, even if the extra information and structure captured today will not be exploited for some time to come.

- Despite its complexity, this solution remained independent of the ARKive problem domain. The generalized editing tool and description structure could be applied to the context of any museum or knowledge collection.

- Ontology tools are already useful for three aspects: structure of the content being created, for a thesaurus of descriptive terms, and for reasoning to infer knowledge about each item from a generalized description of the whole domain. These techniques are still far from mature, and the two areas aren't joining up yet.

- Pokémon aren't Gannets. Building a solution for gannets, which are not fully described before beginning to develop the solution, involves a complex issue, that of nominals.

- Issues of identifying and describing items consistently are significant. This will require either huge effort on maintaining consistency during authoring, or search and access tools that can resolve these gray matches and support validation and quality assurance tasks as part of an overall authoring workflow.

References

Berners-Lee, 1998. T. Berners-Lee (1998) Semantic Web Road map, http://www.w3.org/DesignIssues/Semantic.html

Biosis, 2002. Zoological Record Systematic and Subject Thesaurus http://www.biosis.org/zrdocs/zr_thes/systthes/

Carroll, 2001. Carroll J., (2001) Another RDF Parser http://www.hpl.hp.co.uk/people/jjc/arp/

Cawsey, 2000. Cawsey A. (2000) Presenting tailored resource descriptions: Will XSLT do the job?, WWW9, Amsterdam http://www.cee.hw.ac.uk/~alison/www9/paper.html

Dingley & Shabajee, 2001. Dingley A. & Shabajee P., Use of RDF for Content Re-purposing on the ARKive Project (2001) *International Conference on Advanced Learning Technologies*, Madison

Hunter & Lagoze, 2001. Hunter J. & Lagoze C. (2001) Combining RDF and XML Schemas to Enhance Interoperability Between Metadata Application Profiles, WWW10, Hong Kong, http://archive.dstc.edu.au/RDU/staff/jane-hunter/www10/paper.html

IEEE LTSC, 2001. IEEE LTSC Learning Objects Metadata (LOM), http://ltsc.ieee.org/doc/wg12/LOM_WD6-1_1.doc/

Markey, 1984. Markey K. (1984) Interindexer Consistency Tests: A Literature Review and Report of a Test of Consistncy in Indexing Visual Materials, *Library and Information Science Research*. 6: 155-177

Motta, Buckingham Shum & Domingue, 2000. Ontology-Driven Document Enrichment: Principles, Tools and Applications, *International Journal of Human-Computer Studies* , 2000, 52, (6), 1071-1109.

Wielinga, 2001. Wielinga B.J., Schreiber, A.Th., Wielemaker, J. ,and Sandberg ,J.A.C. (2001) From thesaurus to ontology. *International Conference on Knowledge Capture*. Victoria, Canada

The Virtual Ramp to the Equivalent Experience in the Virtual Museum: Accessibility to Museums on the Web

Liddy Nevile, Motile Research, Australia and Charles McCathieNevile, W3C, France

Abstract

People with disabilities are slowly finding their way on to the Web, up the virtual ramp, as it were. The ramp that is being built is proving popular among many users beyond the community for whom it was first requested. Good practice has ramps integrated into the design of buildings and Web sites from the start, and everyone can feel equally welcome and share in the experiences provided beyond the ramp. The same practices can be adopted on the Web, but often the question is what the best virtual experience that will provide for all inclusively is. The authors argue a range of forms and modalities of resources should be provided to ensure accessibility and richness for all users.

Keywords: Disabilities, Good Practice: Accessibility: W3C Standards; Equivalent Experience; Cognitive Equivalence

Introduction

After a number of years, the communities of people with disabilities are slowly finding their way on to the Web, up the virtual ramp, as it were. The ramp that is being used is, as in many other contexts, popular among many users beyond the community for whom it was first requested. Good practice has ramps integrated into the design of institutions from the start, and everyone can feel equally welcome and share in the experiences provided beyond the ramp.

The authors have been working with a community responsible for a 'museum' that covers several hundred square miles. It contains a collection of 100,000 Quinkan Aboriginal cave paintings and is located in the far north of Queensland in Australia. In this context, the idea of an equivalent experience for those who visit on the Web is of paramount interest. In fact, almost everyone who visits and experiences Quinkan culture and artifacts will be doing so via the Web. Almost all of the 'museum' is inaccessible (the nearest city is 250 miles away) for one reason or another. Currently, work is being undertaken to develop an interpretive centre for this museum and the on-line component of that is of interest here. (See the Queensland Heritage Trails Web site, http://www.heritagetrails.qld.gov.au/fs_quinkan.html)

Virtual tours of Quinkan country will be designed to attract visitors to the real thing, but only a small proportion of visitors will make the journey. The virtual tour will provide the first level of equivalence; the virtual experience for those who cannot have the real experience. The second level of equivalence comes in when the first experience is transformed for the benefit of those who, for whatever reason, do not get access to the first level.

'Equivalence' in the accessibility context means that a user who avoids one presentation of a resource, for one reason or another, can choose an alternative resource to gain an 'equivalent' experience. It is an open question what is equivalent to the range of resources and activities being developed for what is seen as a museum's target audience. So questions for consideration when developing virtual exhibitions of the Quinkan Rock Art include:

- What is equivalent to being shown a 36,000 year old cave painting by an elder of a community that has lived in the region for tens of thousands of years? and
- What is the equivalent of the experience for those who either cannot see, or cannot hear the primary virtual experience?

BEST COPY AVAILABLE

Museums and the Web 2002: Proceedings
Archives & Museum Informatics, 2002, p. 93

For a starting point, the authors present a brief overview of the technology available for accessibility. Then they explore some different approaches to providing access, through the perspective of IMS (http://www.imsproject.org/). They further develop these ideas with particular reference to their similarity with current museum practice, and consider the implications for accessibility. Finally they recommend that museum exhibitors and authors, developing interactive or electronic resources, include planning for accessibility as part of their initial design process.

The Technology of Access

The ramp in the online context is what is defined as universally accessible Web content, according to the W3C Web Content Accessibility Guidelines (http://www.w3.org/WAI/). An effective ramp exists, according to the W3C Web Content Accessibility Initiative working groups' recommendations, when Web content is equally available to all accessibility-standards-compliant devices, whether they are GUI browsers, have mouseless interfaces, or are Braille devices. The guidelines relate to content, authoring tools and access devices and agents.

In the case of a virtual museum, or the online publication and interactions of a museum, such ramps are now required by law in many countries. Currently there is work going on to implement the accessibility requirements but some of this work is not merely technical, and this is especially of interest in the museum context. Generally, the familiar question that is now well-integrated into thinking in the physical world, "How can experiences be provided by museums in ways that provide equal satisfaction for all" is extended to ask, "How can virtual experiences be provided by museums in ways that cater for all needs, including for those who could never enter a physical museum?"

W3C's Web Content Accessibility Guidelines, derived from consultation and consensus among the many and varied disabilities communities, and content and technology developers, call for the provision of equivalent alternative resources and activities for those who cannot use the initial presentation, be it an interactive multimedia object, a video, or just a picture. In this paper, we start to formulate ways of thinking about what this may mean for the design of online experiences.

A video file can have multiple tracks, each catering for modalities that suit different needs. Sometimes such a composite object can be constructed as a unit, with software that handles all the modalities, and on other occasions it needs to be formed by the close association of the sub-objects. It can even be a combination of sub-elements from different sources, as happens when, post-publication of the primary objects, someone else publishes an alternative transformed version, say a caption file that relates to the original resource. Integration of all these files, including the capability to vary the speed with which interactions happen, can be organized by a single integrating language. SMIL (pronounced 'smile' and meaning Synchronised Multimedia Integration Language), (http://www.w3.org/SMIL/) is the recommended XML (http://www.w3.org/XML/) language that can perform this function (http://www.w3.org/AudioVideo/).

Composite objects can be combined in the usual way as Web pages before being published to the Web . In fact, in new work on theories of how authors can learn to create Web resources, the authors have worked with the metaphor of a newspaper. The Web author, according to this theory, is encouraged to develop atomic Web content, aggregate it as accessible composite objects, and then to lay them out on a Web page, to create the Web resource.

An important technology that has made this an appropriate practice is the development of Cascading StyleSheets (http://www.w3.org/Style/CSS/). Using these, and separating the content from its layout is the first step. The content, apart from being composed of objects that are themselves accessible (because they consist of a range of subobjects that themselves cater for a range of needs), needs to be carefully classified or structured. HTML (now deprecated in favor of XHTML) and other XML markup languages can be used to tag objects within the resource so that navigation elements, headings, sub-headings, addresses, etc can be easily picked up by access devices and presented in ways that make them easily identifiable, as they are for those of us who can see them on a web page. CSS is then used to format these objects in ways that make them most useable according to the particular access device being employed. The style sheets 'cascade' in order to allow for a set of them, including a user preferred one, and it is possible for the user to change the layout or presentation by changing the style sheet. In other words, while a GUI device such as a standard browser might use a style sheet that has tags describing how the text, images, links should appear on a screen, a screen-reading device may override the given style sheet and use one that applies different voices as it reads out the differently structured objects.

With the developments made possible by XML, style sheets can themselves be transformed, not just alternated, and there is considerable scope for good devices to make objects and thus resources much more meaningful and useable according to the needs and devices of the user. It is in this context that the authors have been working on their newspaper metaphor for developing accessible resources.

Like a modern museum, the Web is not a static experience of receiving information, but an interactive dynamic space for exploration. Work on ensuring the accessibility of new Web technologies such as Scalable Vector Graphics (http://www.w3.org/SVG/), a way of presenting dynamic animated graphic material in textual form so it can be rendered as images or text, SMIL (as mentioned above) and new interactive building blocks for the Web, is critical for the provision of accessibility. This work forms the base for the virtual ramps.

Where multiple transformations of the same resource are authored by the resource creator, and available to users regardless of their special needs, or access abilities, the creator of the resources is able to communicate directly with the user. Psychologically, as with the ramp into the building, the user has the sense of inclusion and an equal opportunity to participate. The technologies mentioned so far can be used to make available, in appropriate formats and modalities, what has been selected and created by the original resource creator. The point here is that it is not a matter of a creator producing one resource and leaving it to technicians to make alternatives for those with special needs, but rather a situation where the original creator should be encouraged to think *ab initio* of all the different formats and modalities and consider their design part of the main design process.

Direct and Compatible Access

There are two major forms of physical access to online resources (http://www.imsproject.org/accessibility/index.html)

> "A "directly accessible" product is designed so that a person with a disability can operate all on-screen controls and access the content without relying on the aid of an AT [assistive technology]. For example, to be accessible to users with low-vision, directly accessible applications, software, or Web sites offer features to enlarge all controls and on-screen text and are designed with high contrast colors or provide features that allow users to choose appropriate colors. To be accessible to blind users, a directly accessible product should have a keyboard interface with audio output."

A directly accessible Web resource, with a suitable keyboard interface will provide access for many users who are blind, those with physical disabilities and many others who have a temporary disability. Audio output that announces the presence and status of on-screen controls and conveys the atmosphere of the application, software, or Web site assists those with vision disabilities but also the illiterate and the foreign language speakers. A single key method for scanning through choices in the application or software provides access for users dependent on a single switch for input or who are busy using their hands for some other purpose (e.g. knitting or driving).

A "compatibly accessible" application, software, or Web site is designed with AT in mind. This level of access assumes the user has a preferred AT package installed and is relatively competent and comfortable with it. A compatibly accessible product is designed with "hooks" to facilitate ease of use with a screen reader, screen magnifier, or alternative input devices such as adapted keyboards or single switches. These hooks can be implemented by developers ."... Exposing the system cursor, using standard controls and fonts, and following the operating system's human interface guidelines can help make a product compatibly accessible."

Direct and compatible access offers different advantages to different stakeholders in the Web context.

Directly accessible on-line resources have as advantages:

- no requirement for expensive assistive technologies, cutting costs
- reduction of technical complexity
- the opportunity for users to operate any computers anywhere
- direct and designed communication between the creative author and the user, not an AT acting as intermediary, and
- single resources that are suitable for a range of special needs (one size fits all).

Compatibly accessible on-line resources offer:

- consistency of operation across different activities for the user
- lower content development costs
- system-based functionality e.g. an AT package provides text-to-speech capability
- possibly the only access means for some users e.g. Braille, and a single set of programming techniques to be used by the developer for all ATs.

Directly compatible resources allow all users to engage with them 'equally' in the same sense as the ramp into the museum building gives access to all. The experience may be different, because the ability to participate may be different, but everyone has the opportunity to participate in the same activity.

The National Center for Accessible Media (http://ncam.wgbh.org/) has developed a DVD that shows how this can be done with interactive materials that have been made directly accessible for people with a range of disabilities.

Direct access is available when users are provided with transformations of content that suit their access devices.

Imagine an interactive multimedia product that offers a virtual tour of a small area containing some significant rock art paintings in northern Australia. Such a multimedia product would probably contain, in various combinations, text, images, animations, video, sound files, and interactive applications.

Blind persons will not be able to see either the images or the controls to interact with the resources. They will need keyboard replacements for the mouse, for text that describes all images and imagery, and Braille or voice transformations of all text, captions to accompany sound files where these are initially accompanied by images or text, and so on. Deaf persons will be grateful for images and video but possibly not find text easy to read, and may need signing, and this may need to be in their language (e.g. US, or UK standards). A person with cognitive disabilities may need more literal text than others who are not confused by metaphoric representations. Braille readers may need extra time in which to make choices because their transformations are delivered more slowly. And so on.

All of these requirements can be accommodated simultaneously by the creation of composite objects from which the user can choose those required.

An example of a situation when direct access might not be preferred over compatible access is when users of a speech interface are required to learn to use the speech interface associated with the resource instead of being able to rely on their own, and operate with familiar, but compatible access. For blind people who are using a computer every day, it is important to have access to a number of different kinds of software: e-mail, word processing, Web browsing, system tasks, and so on. To deal with this, various systems are available to provide something like an "audio desktop", as familiar to its users as other people are with a visual desktop metaphor of windows placed like overlapping pieces of paper on the screen. For such blind users, it is an interference (sometimes quite literally, stopping their everyday software from working) to have to learn a specially designed system for a unique task when they are used to similar tasks using a familiar interface.

If multiple modalities are available to increase the range of opportunities for direct access, four classes of users may benefit although they would not normally have sought such alternatives. These include users who:

- do not know of their 'disabilities'
- do not consider them worthy of concern
- do not know what to do about them or
- do not know how to, or cannot afford to get assistance.

In addition, as has been noted above, it is not always possible to provide directly accessible transformations of all content so that all users may participate directly.

If a range of resource sub-elements has been developed, so that the users are free to choose among them, in many cases a user with a compatible device will choose to use that in combination with a directly accessible sub-element. In other words, direct access is compatible with compatible access as well-designed direct access sub-elements can be selected by all users, including those who may need to use ATs.

Alternative and Equivalent Access

In practice, however, it may be quite difficult to provide all users with the same resource and achieve the same outcomes. In this sense, it is not the technical difficulties associated with the task but rather the potential to support, or destroy, the original intention of the resource that may be in question. The problem of what to provide for users is not new: in the museum context it has been dealt with extensively as the problem of how to design a good exhibition. An expression often used to describe the quality of good applications that make them useful and learnable to the most users is "low threshold and high ceiling". This means that for users there is a minimal entry level, in terms of technological expertise and domain understanding, and yet the same application can work well for a highly accomplished and knowledgeable person who has strong technical skills.

In the virtual museum context, the question is what will provide a virtual visitor with the richest experience, given that some visitors have special needs. The solution may be one that offers all users an equivalent experience, according to the modality in which they participate. In other words, the provision of resources in multiple modalities may not be sufficient to satisfy the original intention of the resource when the full range of users is taken into account.

Returning to the IMS definitions, (IMS, 2001) we find:

> When considering the accessibility of applications and software for learning, education, and training, it is important to understand the differences between two types of access: equivalent and alternative.
>
> Equivalent Access provides the learner with the same learning activity but it is mediated in a different modality. Providing a course textbook in Braille format, on audio tape, or in digital format are examples of equivalent accessibility. Alternative Access provides the learner with a different learning activity but one that is designed to achieve the same learning objectives. An example of alternative accessibility might be having a mobility-impaired student conduct science experiments in a virtual laboratory, where the same levels of dexterity, strength, and physical access are not required as in a physical laboratory.

What is impressive, and often makes the production of alternative activity worthwhile, is that what was primarily considered as an alternative for people with disabilities, like the ramp, becomes an attraction for an unanticipated community of people. A simple example is provided by the "virtual microscope", developed by the Open University, UK (IMS, 2001) for disabled students but subsequently used by all students because of its ability to achieve key learning objectives more fully than the original exercise involving a real microscope.

It is the cognitive aspect of accessibility that is under consideration here. If the virtual museum experience is designed to replace or be an alternative to a 'real' experience, it will be designed, presumably, to achieve the same effect in the participant's mind. What is not clear, and not answerable, is whether the alternative experience is to provide equal access to the real experience or simply to the virtual experience. Either way, and this will depend upon the circumstances, this will not be easily solved.

Cognitive Equivalence

Many years ago, the authors both worked with computer controlled robot 'turtles' - on screen and off. For a substantial time, teachers saw the Logo 'turtle' as needing to be explained to children in the real world (using a physical object) so they could make sense of the virtual turtle - a triangle of pixels on the screen. In fact, children had no trouble at all identifying with the 'turtle' triangle and also relating to the physical (toy) turtle. There was nothing to suggest that the children did not see these as two separate objects, alternate, and in many ways equivalent, or that they needed to relate to a real turtle in order to be able to relate to a moving image. This thinking led one author through a series of investigations of screen experiences, ultimately concluding that there was nothing 'virtual' about screen objects. In some cases, the virtual object had characteristics that were not possible for a physical object, and so could be used to discover 'artificially', concepts that could not be accessed as easily in the real world.

The most effective use of this artificiality seemed to occur when the virtual or simulated object was 'broken' in significant respects. In the case of the turtle, the screen version did not have to contend with gravity, and could assert its position in co-ordinates. These two features alone made it a most suitable object for programming in a gravity investigation exercise. In the process of attributing gravity to the world of the turtle, children were able to discover important facts about gravity such as that it exerted a constant downwards effect on objects, and that this was not altered by the object's horizontal motion although it changed the path of the object traveling horizontally.

There was 'reality' in the children's minds of equal significance, no matter that it was inspired by a real (physical) or virtual (electronic) object.

If the aim of a museum exhibition is to provide visitors with opportunities to use the exhibition to learn, become more knowledgeable, or whatever it's called, it seems that there is the potential for alternative activities, particularly when supported by computation, to achieve the same goal as the physical exhibition. When considering what alternatives to make available for users with special needs who are accessing a virtual exhibition, the author has the task of deciding what is appropriate to people who have access to only some modalities in which resources can be provided. The authors consider this similar to the familiar task of designing for work with different materials in the real world, and that including the range of modalities extends, rather than alters, the designer's work.

Creating Alternative Experiences

In the end, equivalence is something that is a matter of judgment. The perceptions of authors about what they are trying to convey, and the perceptions of readers (in the broad sense) of what the authors were trying to convey, are not necessarily the same. An equivalent or an alternative can be said to have been a "good" one when readers who used different alternatives come away with the same impressions of what they are being shown.

This is a difficult but mostly soluble problem in a training context where there are specific desired outcomes and measurable descriptions of what the reader should be able to demonstrate having accomplished. But in the more open-ended context of an educational experience, or a museum experience, providing a resource that can be interpreted in different ways is more complicated.

A "minimal equivalence" might be achieved if it is possible for some primary message to be received from any of the available alternatives. For example, a short cartoon like the following might be perceived as conveying the information that text without illustrations does not communicate as well as illustrated text. On the other hand, a "more complete equivalence" could convey something more. In fact, any selection of an alternative will convey a style. Consider the difference between:

> "Supplement text with non-text content."

and

> "Some people can't read writing easily, so add images, sounds, movies and so on to help them understand your message."

and

> "This is a cartoon. On the left is a person looking in confusion at a poster of text. On

> the right the text has had images added to explain it, and "the light has dawned" - the person now understands the poster. The left hand panel has been crossed out indicating it is not the right approach, and the right hand panel has a tick indicating that it is a good approach."

and Figure 1.

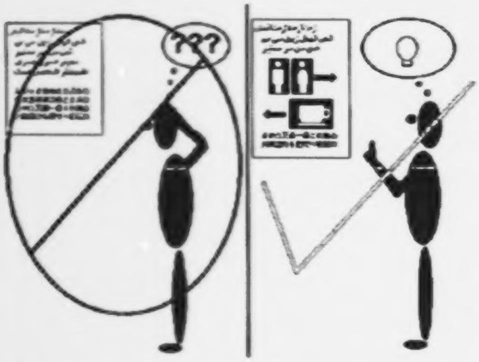

Figure 1: a cartoon http://www.w3.org/2001/08/mmcmn/34a.png

In the work associated with development of artificial environments for learning, the community often refers to such environments as 'microworlds'. One quality of good microworlds is that they satisfy the rule above of having a low threshold but nevertheless a high ceiling. As well, they provide rich opportunities for a range of users without necessarily being able to predict what users will bring to their interaction with the microworld or gain from it. Microworlds usually give control to the users, allowing them to choose what to do in the environment with the designer knowing that whatever is done, it is very likely that something of value will be gained. A physical equivalent of such a microworld might be a sandpit with water for a child.

But one feature of microworlds is outstanding. While they are recognizable for these qualities once developed, it is often not that they have been easy to create, or even that their designers were able to apply a clear set of design practices in order to produce them. Many years of experience with microworlds suggests they are somewhat like other forms of art: not subject to prescription yet enormously rewarding and valuable. The authors believe this to be similar to the work of museum exhibition design, and feel confident that museums will embrace accessibility and demonstrate their expertise in a slightly more complex world with the usual high standards of success.

Conclusion

Having considered the tools with which museums might work to build the virtual ramp, and the growing demonstration of talent within museums to create on-line resources of all kinds, the authors are optimistic that the museum community will be able to make a substantial contribution to the more general effort to make the Web a rich, exciting and educational environment for all.

A young Aboriginal man recently spoke of his efforts to make sure the rock path he was building up to a rock painting site was appropriate for the site (Steffensen, 2000). We believe the ramps that give access to virtual visits to the Quinkan Rock Art of far north Australia can be just as inclusive and rewarding.

References

Steffensen, V. presentation (without published paper) at International Rock Art Conference, hosted by AURA, Alice Springs, Australia: June 2000.

BEST COPY AVAILABLE

Adding Value to Large Multimedia Collections Through Annotation Technologies and Tools: Serving Communities of Interest

Paul Shabajee, Graduate School of Education and the Institute for Learning and Research Technology (ILRT), University of Bristol, Libby Miller, Institute for Learning and Research Technology (ILRT), University of Bristol, and Andy Dingley, Codesmiths, UK

Abstract

A group of research projects based at HP-Labs Bristol, the University of Bristol and ARKive (a new large multimedia database project focused on the world's biodiversity based in the UK) are working to develop a flexible model for the indexing of multimedia collections that allows users to 'annotate' content utilizing extensible controlled vocabularies. As part of the educationally focused ARKive-ERA project, a series of models for user 'annotation' have been developed. The need for these types of user support and tools was identified while conducting pre-design user studies with specialist user groups. The needs center around the limitations of current on-line museum and library systems that do not provide support for users to annotate or 'tag' multimedia objects of relevance to their particular 'community of interest' or with specialized indexing terms. Tagging would enable specialized resource discovery and knowledge sharing with other members of their communities. One example is that of University Lecturers and Researchers studying a particular type of animal behavior. They may wish to identify all relevant images or video of that particular behavior and annotate them as good illustrations of aspects of that behavior. However, significant issues arise over, for example, the validation of information, access control and the use of such annotations by the resource discovery tools. The paper explores these and other issues and problems involved, and explains how the various models can help provide solutions to key problems and thus meet the needs of a diverse range of 'communities of interest', thereby adding significant value to on-line multimedia collections.

Key Words: community annotation, flexible publishing, semantic web, ontologies, collaboration

Introduction

The ARKive-ERA project is focused on investigating how best to design the underlying technological infrastructures to enable large multimedia database systems to maximize the educational potential of their multimedia assets, for users from very diverse range of backgrounds and in a wide variety of contexts. The focus for the research has been the ARKive project (http://www.wildscreen.org.uk/arkive/), a large multimedia Web-based database system under development, containing diverse data related to endangered animal, plant and fungi species and their habitats as well as more common UK species.

ARKive is characteristic of many large digitization projects; during its *initial phase* of development it will contain data profiling some 2000 species and their habitats. This will take the form of approximately 9,000 minutes of digitized video and 30,000 still images along with hours of audio, maps, textual information and other supporting media and educational materials. These assets are donated by a diverse range of commercial and non-profit organizations as well as by individuals.

Essentially ARKive is a 'community project' insofar as it is part of and relies on a community of organizations and individuals who have an interest in sharing access to rich multimedia resources focused on biodiversity.

ARKive type projects are designed to serve the needs of their diverse potential users by providing tools for individuals and communities of users to 'annotate' the content of the database so as to make the content more valuable to others with similar interests.

It is important to note that we define *annotation* as metadata (see below) created after the creation of

BEST COPY AVAILABLE

Museums and the Web 2002: Proceedings
Archives & Museum Informatics, 2002, p. 101

the content. It is this post hoc nature ("a note added to anything written", Oxford English Dictionary, 1998) that represents a considerable expansion of its usefulness, as a means of adding value to content, because it now allows people other than the original content author to add metadata descriptions.

The Challenge

As part of early work to identify the key requirements of ARKive and similar projects, it became clear that the diverse range of potential users includes school children and their teachers, media researchers, conservation scientists, customs officers, university lecturers and students and the very many people with personal rather than professional or educational interests in wildlife, to name but a few.

Each of the groups listed above is itself diverse with respect to the particular needs or desires of ARKive or similar systems.

We conducted a small-scale interview survey with University Lecturers about their likely uses and needs of ARKive with respect to supporting their teaching activities. As a result we identified that a key need of this 'group of users' was to be able to search for multimedia resources to 'illustrate' concepts when presenting to and supporting their students. Specific examples included infanticide, drug induced behavior, life strategies of plants, inter-specific competition, identifying and classifying organisms, tropic levels, echolocation, binocular vision, and harvesting theory. Lecturers from different sub-domains (e.g. behavioral biology and ecology) suggested different terms.

This small example shows that even within sub-groups of one relatively well-defined group of users the 'resource discovery needs' alone were complex. Indeed it quickly became clear that all of these sub-groups or 'communities of interest' have their own specialized vocabularies and concepts that they would like the resources indexed under so as to support their resource discovery needs.

Not only did they want to be able to search for assets using specialist vocabularies, but they also wanted ideally be able to find 'good' examples of assets that illustrated a particular aspect of a particular concept.

These are not unreasonable requests as clearly it is not practical to browse 9,000 minutes of video, or even a small sub-set, in the hope of finding a good example to illustrate aspects of, for example, 'harvesting theory'. However for ARKive it is simply not feasible to index every asset with the terms relevant to all possible communities of interest.

The problems can be expressed more clearly by using basic concepts from Information Retrieval (IR) literature (e.g. Chowdhury 1999).

- **Precision** (number of relevant documents in results divided by the number of retrieved objects)
- **Recall** (number of retrieved documents divided by the number of relevant objects in the collection)
- **Specificity** (level of detail of indexing of an object i.e. how fine grained is the indexing)
- **Exhaustivity** (completeness of indexing of object using available vocabulary)

If we use the example above, the University lectures (not unreasonably) want to use their specialist vocabularies to search ARKive's multimedia archive and have high precision and recall from the system based on those terms and queries constructed from them.

When actually doing the indexing, ARKive has to balance the levels of specificity and exhaustiveness of their indexing to make the task tractable within the limits of available resources and time.

It is useful here to refer to some well-defined 'specialist vocabularies' to get some insight into the scale of the challenge.

- The Biosys, Zoological Record http://www.biosis.org.uk/products_services/zoorecord.html is indexed using an extensive thesaurus of around 10,000 terms developed over more than 20 years.
- The GEneral Multilingual Environmental Thesaurus (GEMET) was developed by the European Environment Agency (EEA) together with a co-operation of international experts to serve the

needs of environmental information systems. Analysis and evaluation work produced a core terminology of 5,400 generalized environmental terms and their definitions. (European Environment Agency, 2002)

- Medical Subject Headings (MeSH) http://www.nlm.nih.gov/mesh/meshhome.html contain more than 19,000 main headings. There are also 103,500 headings called Supplementary Concept Records within a separate chemical thesaurus.
- The Art & Architecture Thesaurus (AAT - http://www.getty.edu/research/tools/vocabulary/aat) is maintained by the Getty Research Institute (see http://www.getty.edu/gri/) and the thesaurus keeps growing. The AAT contains about 120,000 terms covering objects, textual materials, images, architecture and material culture from antiquity to the present. (J. Paul Getty Trust, 2002)
- UK National Curriculum (for schools) Metadata Schema contains some 2000 'subject keywords' (Qualifications and Curriculum Authority, 2002)

While any particular collection of multimedia will not necessarily contain objects which are appropriately indexed under many of the terms from each and every available specialist vocabulary, it is none-the-less possible that some terms from all of these will be applicable to some objects.

These issues are relevant to many other types of projects, not least those involved in the development of cross-searching of multiple databases which have been indexed using different indexing schemas (e.g. Clark, 2001). Much work is going on with respect to providing ways of robustly mapping between different schemas. These issues are discussed again below.

Indexing, Metadata, Interoperability and Ontologies

It is beyond the scope of this paper to review the extensive literature on the indexing and applications of 'metadata' (see below) to multimedia objects, the issues of interoperability and related technologies. See Gill and Miller (2002) for an overview of the key issues with regard to 'digital cultural content', and below for some examples of relevant projects. However a brief overview is necessary and useful in providing additional background to the remainder of the paper.

Metadata is broadly defined as 'data about data' (Gilliland-Swetland, 1998). The traditional library catalogue index card is a classic example of metadata. The publication date, author, title, publisher, dewey decimal code… are *'metadata elements'* within a clearly defined *metadata schema* and *scheme* (list of metadata elements, allowed states of those and relationships between them).

As can be seen from this example, there are different types of metadata. Gilliland-Swetland (1998) distinguishes between 5 types:

- **Administrative:** Metadata used in managing and administering information resources
- **Descriptive:** Metadata used to describe or identify information resources
- **Preservation:** Metadata related to the preservation management of information resources
- **Technical:** Metadata related to how a system functions or metadata behave
- **Use:** Metadata related to the level and type of use of information resources

Descriptive metadata is of most relevance to the challenges outlined above, but in principle the issues apply to all the types.

The interoperability of metadata i.e. the ability of different information systems to inter-operate or be compatible with each other's vocabularies is seen as a fundamentally important issue in the development of Web-based information systems (Gill and Miller, 2002). This is because it is valuable if two (or more) systems holding data on similar things can be reliably cross searched and/or share data. Many standards, initiatives and projects are in place to develop systems that will be able to interoperate at a vocabulary and semantic (meaning) level (e.g. W3C 2002b, Miller 2001, see also below).

Part of the development of interoperable Web-based systems includes the creation of systems that utilize semantically interoperable ways of describing things, characteristics of things, and the relationships between them. These 'ontologies' (e.g. Ontologies W3C initiative, W3C 2002b) take the form of structured machine readable representation of the knowledge.

> Just like people need to have agreement on the meanings of the words they employ in their communication, computers need mechanisms for agreeing on the meanings of terms in order to communicate effectively. Formal descriptions of terms in a certain area (shopping or manufacturing, for example) are called ontologies and are a necessary part of the Semantic Web. RDF [Resource Description Framework], ontologies, and the representation of meaning so that computers can help people do work are all topics of the Semantic Web Activity.*(W3C 2001b)*

These developments form what can be seen as part of a larger movement in Web technology development towards a more semantically interoperable Web (W3C 2001a, Berners-Lee et al 2001) in which information is globally interoperable.

There are significant difficulties with building ontologies and applying them to bodies of information. Ontology creation and application is a very specialized and time-consuming activity. Even more difficult is mapping between ontologies, especially those written by different communities of interest. An ontology provides a machine processable hierarchy of terms, but not all of the intentions of the ontology creator are encoded into the description of the ontology. Therefore mapping between them is prone to errors of interpretation.

Part of a Solution: Community of Interest/Expertise Annotation

The development of more semantically interoperable Web-based technologies seems to promise the ability to solve part of the challenge outlined above; namely, that of enabling machines (computers) to relate terms from different specialist vocabularies about what are essentially the same thing or concept and thus being able to map existing terms to the specialist vocabularies (including other languages). However they do not provide a solution to the problem that members of specialist interest groups will want to describe (apply metadata to) data in ways relating to totally different concepts.

A simple example makes the issue clearer:

> Imagine that there is a database of 100,000 images of people in a wide variety of different settings, say developed for a news agency. The database may be indexed using terms relating to identification of people (name, age,...) and event (time, place...) as well as administrative, preservation and technical metadata. This is because those are the important characteristics to those who originally setup the database.
>
> Now milliners might see great potential for studying how people use hats. The database is likely to be a very useful resource:, they could search to see how the style of hats has changed over time, or what types of hats are most popular; they could answer many more specific questions, e.g. what percentage of women have bows on their hats? or wear a particular type hat at a particular time of year?... however the database is not indexed using the concept of 'hat' and so it is not possible to interrogate it to find the answers to these questions.
>
> Imagine now that someone else, a landscape architect, comes across the database and sees that it could be used to study how public seating is used in urban settings...
>
> In each of these examples the collection *could be* of very great value to the user, but the existing indexing was not originally designed with these uses in mind and so it is not. It is in these cases that *community annotation* of a collection, could offer the key to meeting these needs and thus greatly extend the scope and value of an on-line collection.

In the example above the individuals are from communities of 'milliners' and 'landscape architects'. They could annotate the images with specialist indexing terms used by their communities, ideally from ontologies developed to facilitate a semantically consistent representation. However it might be that they simply want to add 'notes' to particular images for others from their communities to find or make a hyper-link to a page of more detailed complementary information or case studies of that kind of example.

Models of Community Annotation

This section outlines a number of models of community annotation that we have identified, and that we believe help meet the diverse needs of different 'user communities' and the database system developers such as ARKive.

Figure 1 shows via a much simplified Venn diagram some of the communities of interest of an ARKive type project. There are the more traditional 'target users', in ARKive's case, those with interests in biodiversity and wildlife media and their sub-communities A1, A2, A3… (e.g. different sub-disciplines, phases of education…), and ARKive's own staff. However there are other 'communities of interest' (B, C, D and E) that lie outside or may have some degree of overlap with the original target community or communities.

Whilst the discussion here is focused on the use of community annotation to apply 'specialist' indexing to objects, the majority of the issues discussed are the same for other kinds of annotation. Examples include case studies in the use of a particular image or media type, or notes relating to the object (e.g. what it shows, interesting facts, controversies).

Possibly the most critical issues related to 'community annotation' for the organizations behind ARKive-type Web sites are related to the quality, accuracy and relevance of any annotation. One of ARKive's fundamental values is ensuring that it provides scientifically accurate and up-to-date information.

If users are given tools that enable them to add/link 'metadata' to multimedia assets, many issues arise about how that annotation can, or should be, made available to other users of the system. Below we have outlined four models of annotation that we have developed to illustrate the issues.

1. **Trusted members of trusted communities**: ARKive already has a group of 'trusted' experts and organizations that provide and validate information which ARKive collates and publishes on its Web site. This approach would expand this single group to wider communities; for example, a 'trusted organization' could provide a list of potential members that it states are competent to annotate ARKive resources. Likely the resources and 'metadata terms' that they are permitted (by ARKive) to annotate and use, would be limited to a specific set within their area of competency and that the types of annotation or vocabulary of concepts would also be limited (see consistency below).

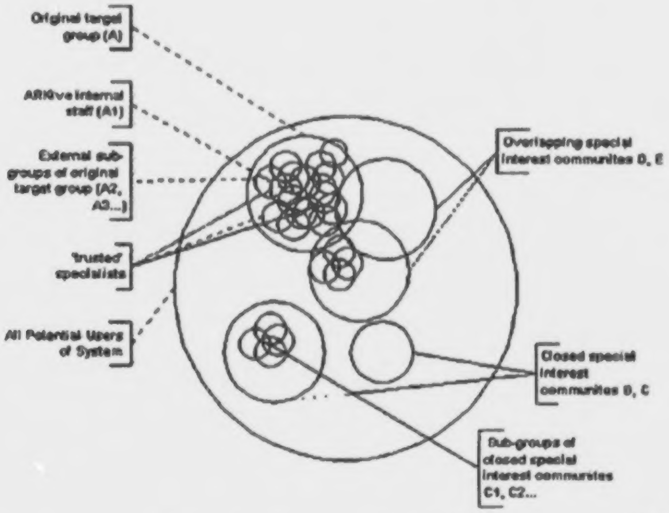
Original target group (A)
ARKive internal staff (A1)
External sub-groups of original target group (A2, A3...)
'trusted' specialists
All Potential Users of System
Overlapping special interest communites D, E
Closed special interest communites B, C
Sub-groups of closed special interest communites C1, C2...

Figure 1 – Communities of Interest

BEST COPY AVAILABLE

These 'trusted members' could be given usernames and passwords and on-line tools to annotate ARKive resources. These annotations would thus in principle be 'ARKive approved annotations' and thus in effect integral to the ARKive cataloguing and indexing system. However, it may be that only particular communities of users get access to particular sets of annotations.

2. **Self-selecting communities**: Much as there are self-selecting, open and closed discussion groups on academic mail lists such as JISCMail (http://www.jiscmail.ac.uk), it would be possible to provide annotation tools for a self-selecting and administering community of users.

   Such communities normally have 'community leaders' who administer them on a day-to-day basis. They vary greatly in nature from highly structured and constituted to loose with similar interests.

   It could be that only members of a 'closed' community get access to the annotations added by their group, or that these are made available to generic ARKive users but with a 'disclaimer'.

3. **Open annotation:** in this model any ARKive user could annotate an object. This is similar to ratings systems which exist on e-commerce sites such as Amazon.com (www.amazon.com). The degree of 'openness' may range from totally 'open' to a more 'mediated' approach in which some validation or quality criteria are applied (see 'Consistency and Quality Control' below).

   In this case, all annotation would have to have some form of 'disclaimer' as ARKive would have relatively little control over the quality, accuracy or relevance of any annotation.

4. **Third Party Annotation:** ARKive users might want to produce 'third party' sites to **draw** together resources from ARKive and other sites; the simplest example of this would be a list of links to other Web sites. However, these third parties might also produce their own 'annotations' for ARKive resources; e.g. they might use a very subject-specific vocabulary. 'Annotating' resources from diverse sources has the advantage of enabling resource discovery across multiple information sources but with a particular specialist vocabulary and personal control.

   ARKive would not have control over this form of annotation but might want to provide infrastructures and tools to support such annotation.

Existing Projects

Many Web-based projects and sites use some form of annotation to 'add value' to their data. Each example below utilizes one of the approaches above:

- Amazon.co.uk http://www.amazon.co.uk provides customers with the ability to add comments about a product and give it a star rating. This gives future users the 'added value' of hearing the views of others who had read, listened to, watched etc... the product. All users can see all annotations once they have been vetted for compliance with Amazon's guidelines. For guidelines see http://www.amazon.co.uk/exec/obidos/subst/misc/author-review-guidelines.html/026-1836225-8318031)

- PseudoCAP::Pseudomonas aeruginosa Community Annotation Project (http://www.cmdr.ubc.ca/bobh/PAAP.htm) allows Web-based annotation tools a to be used by members of a closed community. Since they aim toto:

- "... improve the quality of analysis of the *Pseudomonas aeruginosa* PAO1 genomic sequence, and to ensure the development and widespread availability of genetic tools to analyze Pseudomonas gene function, this *Pseudomonas aeruginosa* community annotation project (PseudoCAP) was initiated to enlist the expertise of volunteer Pseudomonas scientists in annotating the genome sequence. Annotations provided by the Pseudomonas scientists were subjected to peer review and used to aid the final genome annotation that was published. All participants in this project for the publication of the genome sequence have been acknowledged in the genome paper..."

- Slashdot (http://slashdot.org) is an example of a 'community news portal.' It allows users to upload news and then others can take part in a (threaded) discussion based on the original submission.

110

- Gimp-Savvy.com (http://www.gimp-savvy.com) is a Community-Indexed Photo Archive which provides simple tools for users to add indexing terms to images in an on-line image database (http://gimp-savvy.com/PHOTO-ARCHIVE/)
- Berkman Center for Internet and Society (http://eon.law.harvard.edu/cite/annotate.cgi) has developed an on-line tool (Annotation Engine) for users to annotate on-line documents by placing a link in the text at the point that the user wishes to annotate
- FishBase (http://www.fishbase.org/) is an interesting example of a specialist site which is run and managed as a community. FishBase contributors have passworded access to edit and add data to the very extensive underlying database of scientific data and multimedia resources.

More generally the W3C are looking to develop 'annotation standards' under the Annotea project (W3C, 2002a) to allow users to collaboratively annotate Web pages. In parallel with these developments, the standards which support the use of metadata descriptions are expanding; e.g. the MPEG-7 standard (Martínez, 2001) for the 'content description' of multimedia includes a comprehensive 'Description Definition Language' which allows complex description of multimedia objects.

These projects all provide tools that enable users of different types to add value to the 'collections' by adding annotation.

Implementation of Community Annotation – Issues

Use of and Access to Annotation data

As can be seen above, the use of and access to any annotation is an issue inseparable from that of the under lying model of annotation. Probably the most fundamental issues are deciding who has access to any annotations and how any annotations (explicit, e.g. case studies, or implicit, e.g. search terms) are used and their use signaled. Some approaches follow.

1. **'Trusted community annotation'** might effectively be transparent to any users and simply be utilized as core ARKive tagging or alternatively might be provided via a specialist search engine option or provided with a disclaimer.
2. **'Self-selecting community annotation'** might be available to all users via a disclaimer or only available to 'subscribers' to that annotation (i.e. they explicitly request access) or only to members of the community.
3. **'Open annotation'** could be 'transparent' or have access controlled as in the previous two cases.
4. **'Third party annotation'** would mean ARKive would have no control over access to or use of this type of annotation. For user groups this might be seen as an advantage as the system is 'independent'.

Consistency and 'Quality Control'

Another fundamental issue with respect to using community annotation to assist with 'indexing' metadata is ensuring that there is consistency in both the terms used and the application of those terms, which relate back to precision and recall (see above).

The main solution to the issue of consistency has historically been to utilize a 'controlled vocabulary' with clear instructions about what those terms relate to; e.g. library catalogue systems. In particular, controlled vocabularies appear to improve consistency where indexing is being conducted by a number of 'indexers' (Markey, 1984). In the case of ARKive there is already a controlled (bespoke) vocabulary for a number of aspects of the data and metadata used to describe the multimedia objects.

In order to annotate objects with relevant terms/concepts, it seems necessary to provide not only a controlled vocabulary but also a highly structured conceptual framework on which the vocabulary is based. This is because of the very large range of concepts that are covered by ARKive content, including bio- and bio-geographic sciences, wildlife film making, conservation and sustainable development, and educational uses.

These are broadly quality control issues; key questions for any system will relate to the degree of 'quality control' required for any annotations. This may extend from a check that annotations are not obscene or contravene legal requirements (e.g. libel laws) all the way to full multilevel verification by appropriate experts with formal 'sign off' of any new annotations.

The degree to which this is appropriate must depend on the nature of the system; e.g. a 'closed community' in which members of the community are the *only* users who can access the annotation may require no formal quality control (other than legal issues) from the collection/Web site owners. However 'trusted community' annotation that is accessible to all users may require significant quality control.

Tracking Annotation & Access Control – Annotation Metadata

If community annotation is used, there must be systems in place to manage the new data. That is the ability to track and maintain the annotations; it is necessary to have metadata about the annotations. Cross et al (2001) show how this kind of data can be created and maintained. The potential value of this data is significant, as it potentially enables users (internal to the organization or external) to query the system to say "show me all the annotation to the collection (or subset) made by person 'x' or members of community 'y'". This forms the basis of providing controlled access to the annotation data.

Extensibility

A further requirement of any system of annotation focused on adding value to a collection by 'indexing' is that it be extensible; i.e. that new terms can be added in a coherent and meaningful way.

For example when a new 'term' is added it must be done in such a way that concepts of which it is a sub-element (e.g. pecking might be a sub-set of feeding or defensive behaviors) retain conceptual integrity, e.g. the term 'pecking' should not be applied to an object that does not have a related parent concept (e.g. feeding) or that existing parent concept must (henceforth) be made to apply to the object as well. Hence there may be a need to create new non-overlapping sub-categories of pecking; e.g. feeding:pecking and defensive-behavior:pecking.

This simple example shows that the creation of any conceptual representation will be very problematic. However the current authors believe that without such a framework, effective use of annotation would be very problematic if not impossible to manage and monitor. Heflin and Hendler (2000) explore the complexities of making changes to formal ontologies and some the many associated problems.

Other issues include how to deal with and represent 'controversial knowledge', 'fallacies', 'old knowledge' and other forms of 'inconstancies' in any knowledge base. There are no simple answers to these problems. Once again the most appropriate solution will depend on the particular situation; e.g. in the case of relatively open annotation such as gimp-savvy.com http://www.gimp-savvy.com it may be appropriate for any user to be able to add indexing terms (given legal considerations are dealt with, see above) whilst in the case of 'trusted community' annotation, changes to the ontology or vocabulary used might require a formal meeting of some form of 'expert panel'.

However whatever form it takes, we argue that there must be some system(s) to facilitate such extension of the available terms and concepts if the overall systems are to be effective and sustainable.

A further fundamental problem is that community annotation using extensible annotation vocabularies and schemes is post-hoc and thus, unless every object is systematically annotated, it is very likely that the some objects will not be tagged with a new type of annotation e.g. a particular indexing term, when it 'should' be. Thus the annotation or indexing becomes inconsistent across the collection. There seems to be no simple solution to this problem other than systematic annotation. However as outlined in the next section, the use of 'semantically aware' tools may provide a means of optimizing the completeness of annotations across a collection where (as in most cases) there are time and resource constraints.

Semantic Bootstrapping

One very interesting requirement that we have identified for all of the models described above is 'semantic bootstrapping'; that is, when a collection has been 'indexed' from one perspective (i.e. for its primary target use or user group) using one set of vocabularies, it is necessary to have or create some kind of 'semantic hook(s)' in the data to allow users to begin the process of indexing the collection from the new perspective, using the new vocabulary.

This is a form of 'semantic bootstrapping', conceptually related to the idea developed by Pinker (1984) to refer to his postulated process by which children 'semantically bootstrap' or learn syntax from some form(s) of built in 'semantic categories' and contexts.

Ontology-based tools (see above) could allow existing ontologies to be linked to the already-present vocabularies or ontologies via concepts common to both existing and new domains.

Concept extraction tools such as the 'Non Zero Match' tool were developed at the University of Bristol (http://nzm.dig.bris.ac.uk/index.html). The tool allows users to auto-index text-based documents using concepts defined by a list of words/ phrases with positive and negative weights. E.g. say a 'car' by defining the concept via the occurrence of a set of words or phrases 'registration number', 'steering wheel', 'make', 'model' etc... the parser then processes the whole corpus of documents indexing the documents under the appropriate concepts. Thus by using existing text or indexing/markup it would be possible to create new concepts to help 'bootstrap' the new indexing.

Another example is described by Bobrovnikoff (2000) using the DIPRE (Dual Interative Pattern Relation Extraction) algorithm, to recognize pattern in existing data. Auto-indexing of still and moving images also provides the potential to extract and index new concepts; e.g. in the example above of looking for 'hats' in the database of images of people. See Campbell et al (1997) and Lew (2000) for examples of this approach.

There are various forms that semantic bootstrapping could take, with various levels of automation. It could be a time-consuming and highly skilled manual task, effectively re-indexing the database manually by re-cataloguing by placing the images within an ontology used by a community of interest, or by using a controlled vocabulary. At the opposite end of the spectrum, the images could be auto-classified using specialist tools for pattern recognition.

Somewhere in between is a stored search by a subject expert. For example, when people search Google for a particular topic, they use their knowledge of their subject area and their common sense coupled with their experience of the content of the Google database itself to choose search terms that will accurately retrieve the information they require. For example someone looking for 'flying things' would use more specific search terms like 'bird', 'helicopter' 'parrot'.

An annotated stored query of a database by someone with knowledge of the indexing terms used in the database and the specialist subject knowledge of the community of interest would enhance the value of the database to that community. Such an annotation would provide fast approximate information for that community. An example might be a search of a database for photos of people dressed for a 'formal event' to get pictures of hats.

If there were time, the user could go through the images found in this way to check if the retrieved pictures were in fact pictures with hats, discarding those that were not. However, even a quick annotated search could provide added value.

Some form of 'semantic bootstrapping' will be essential in making annotation work effectively for communities. Different types of 'semantic bootstrapping' tools are likely to assist with different types of problem, and hence it is likely that what is needed is a suite of tools rather than one single tool or approach.

Moving Forward

We are working to formalize the models outlined above and are developing more detailed technical requirements for the implementation of the models. The ideal is that we design an approach that can allow ARKive type organizations to implement any and all of the models of annotation outlined above.

In parallel we are investigating the advantages and disadvantages of the different models in different contexts in order to help developers make decisions about which ones are the most appropriate for their particular needs and contexts. Parallel investigation into different approaches to 'semantic bootstrapping' and development of appropriate tool sets will continue.

Conclusions

Community annotation offers the developers of large multimedia database systems the ability to support specialist 'communities of interest' and thus enhance the value of their data. There are many technologies available and under development that would support this approach; some projects are already utilizing them.

This paper has dealt primarily with the annotation of multimedia objects with specialist indexing/resource discovery terms and the associated technologies; however, the issues are similar for more generic types of annotation.

The four models of community annotation outlined in the paper provide a framework for the development of community-based approaches to enhance the value of Web-based museum and multimedia collections for specialist communities of interest.

There are many implementation issues that remain highly problematic, in particular the coherent and consistent extensibility of vocabularies and the development of 'semantic bootstrapping' tools.

However, it will likely be possible, in the short to medium term, to find solutions by assessing needs and matching solutions in each specific case. In the longer term, we hope that the on-going development of a more 'semantically interoperable' Web and associated technologies will lead to the creation of sets of approaches and tools to make the implementation of community-based annotation relatively simple and effective.

Acknowledgements

The ARKive-ERA project is funded by HP-Labs, Bristol. The authors wish to thank Dave Reynolds of HP Labs, for discussing, exploring and expanding the ideas related in this paper.

References

Berners-Lee, T., Hendler, J. and Lassila, O. (2001) The Semantic Web, Scientific American, May 2001.

Buckingham Shum, S., E. Motta, et al. (2000). *International Journal on Digital Libraries* **3**(3): 237-248.

Chowdhury, G. G. (1999) Introduction to Modern Information Retrieval, Library Association Publishing, London.

Clark, J. (2001). "Subject portals." Ariadne(29). Available on-line: http://www.ariadne.ac.uk/issue29/clark/

Cross, P., Miller, L., and Palmer, S. (2001). Using RDF to Annotate the (Semantic) Web. K-CAP Workshop Knowledge Markup & Semantic Annotation, Victoria B.C., Canada.

DELESE (2001). Digital Library for Earth System Education (DELESE) Web site. http://www.dlese.org/

Eakins, J. and M. Graham (1999). Content-based Image Retrieval, JTAP (Joint Technology Applications Programme).

European Environment Agency (2002). GEneral Multilingual Environmental Thesaurus (GEMET): The GEMET 2.0 Approach. Available on-line: http://www.mu.niedersachsen.de/cds/etc-cds_neu/library/Gemet.pdf

Gill, T. and P. Miller (2002). "Re-inventing the Wheel? Standards, Interoperability and Digital Cultural Content." D-Lib Magazine 8(1).

Gilliland-Swetland, Anne J. "Setting the Stage: Defining Metadata" in Introduction to Metadata: Pathways to Digital Information, Murtha Baca, ed. (Los Angeles: Getty Information Institute, 1998) Available on-line: http://www.getty.edu/research/institute/standards/intrometadata/2_articles/index.html

Heflin, J. and J. Hendler (2000). Dynamic ontologies on the Web. Seventeenth National Conference on Artificial Intelligence (AAAI-2000).

Lew, Michael (2000) Next-Generation Web Searches for Visual Content, IEEE Computer, 33(11) p46-53, November, 2000

Markey, K. (1984) "Interindexer Consistency Tests: A Literature Review and Report of a Test of Consistency in Indexing Visual Materials.", Library and Information Science Research, 6, 155-177.

Martínez, J. M. (2001). Overview of the MPEG-7 Standard (version 6.0), MPEG (Moving Picture Experts Group). Available on-line http://mpeg.telecomitalialab.com/standards/mpeg-7/mpeg-7.htm

Miller, P. (2001) Interoperability Focus homepage. Available on-line http://www.ukoln.ac.uk/interop-focus/.

Campbell, N.W. Mackeown, W. P. J., Thomas, B. T, and Troscianko, T. Interpreting Image Databases by Region Classification, Pattern Recognition (Special Edition on Image Databases), 30(4):555-563, April 1997.

Qualifications and Curriculum Authority (2002). National Curriculum Online Metadata Standard Overview, Available on-line http://www.nc.uk.net/metadata

Schreiber, A. T., B. Dubbeldam, et al. (2001). "Ontology-Based Photo Annotation." IEEE Inteligent Systems May/June 2001: 2-10.

W3C. (2001a) Semantic Web. Available online: http://www.w3.org/2001/sw/"

W3C (2001b), XML-in-10-points, Available online: http://www.w3.org/XML/1999/XML-in-10-points

W3C. (2002a) Annotea Project Homepage. W3C. Available: http://www.w3.org/2001/Annotea/.

W3C. (2002b) Web-Ontology (WebOnt) Working Group Homepage. W3C. Available: http://www.w3.org/2001/sw/WebOnt/.

PGP
"Pretty Good Practice"

Content Management for a Content-Rich Website

Nik Honeysett, J. Paul Getty Trust, Los Angeles, USA

http://www.getty.edu

Abstract

Over the last year the J. Paul Getty Trust's Web presence has evolved from a group of disparate, independently maintained Web sites into a homogenous consistently branded one – getty.edu. This transformation recently culminated with the implementation of a leading Content Management System (CMS). There were and are many process-changes and challenges in implementing a CMS in an institution such as the Getty. These issues are not unique and span the gamut from social, to business, to technical. This paper will highlight the major issues, describe the route the Getty took, and give an insight into the functionality and capability of a leading CMS application for a content-rich museum Web site.

Keywords: Content Management, Templates, Workflow, Virtualisation, XML

Introduction

The Getty is a campus of six programs: Museum, Research Institute, Conservation Institute, Trust Administration, Grant Program and Leadership Institute. All these programs contribute content for the Web site, resulting in a very content-rich environment. It currently includes details on 4,000 works of art; 1,500 artist biographies; 3 hours of streaming video; past and present exhibitions; art historical research papers and tools; conservation research papers and tools; on-line library research tools; detailed visitor information; an event calendar, and a parking reservation system.

Up until a year ago, each program independently maintained its own portion of getty.edu. A gateway homepage was placed at the top level, but as soon as one drilled down into the Website, it was immediately apparent that there was no consistent design, navigation, or look and feel, and no way to navigate around the site without going back to the homepage.

A trustee-level decision was made to present the Web site with a consistent design and look and feel – a Web group was formed and charged with this task. Accomplishing this task would include the implementation of a Content Management System (CMS) to ensure that all the programs could continue to contribute content to the Web site but in a managed and decentralized way. Rather than implement the redesign and CMS simultaneously, the redesign was implemented first, but significant effort was assigned to preparing the HTML pages for later CMS integration.

The Web site was re-launched in February, 2001, with a unified design and functionality and a more thematic treatment of the content. The task had required that the webgroup temporarily 'own' all the content and the processes associated with getting that content to the Web. The knowledge gained during this process was crucial in being able to identify the real requirements of any CMS that the Getty would need to implement.

Maintenance Burden

There was a significant burden to re-launching the Web site before implementing a CMS. To guarantee that ongoing development and maintenance on the Web site adhered to the new rules and style guide, all content still had to funnel through the webgroup – the new design had committed us to a frequent subsite update regime. The GartnerGroup, whom the Getty use to benchmark themselves against similar organizations and business processes, publish a graph which exactly summed up where we stood in the Web life-cycle. Even with a webgroup of 18 people, there were few resources to develop new content.

Preparing For a CMS

To prepare the Getty community for a CMS, which might take 6-9 months to select, the webgroup temporarily established its own managed environment to create and deploy content to coincide with the launch of the redesign. On the infrastructure side, we implemented a three-tier staging environment

BEST COPY AVAILABLE

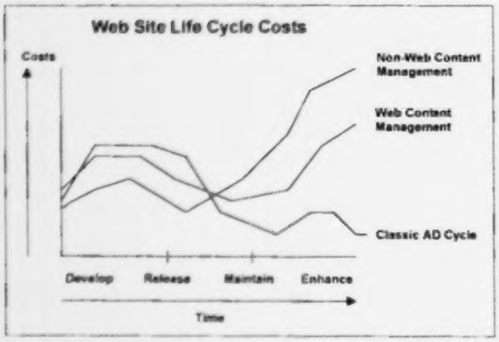

Fig 1: Website Life Cycle Costs. © GartnerGroup 1999

(see figure 2) and a custom-written application to 'push' content between tiers. This was a fairly high-level tool operating on directories rather than individual files. On the business side, we established a strict push schedule, which meant that content could only be moved into production twice a week – although this could be fast-tracked for emergencies on a case by case basis.

Initially there was great resistance from the programs to this regime. The 'anarchic' nature of content generation and deployment that had established itself at the Getty meant that this was a considerable culture shock. The programs perceived a 'loss of control' of their content, and objected to the delay between pushes. Fortunately, the reality was different: the schedule of pushing every other day allowed for only one day to create and review content before it went live. The programs soon fell into the regime and altered their business processes to accommodate the schedule. With the CMS implementation we could afford to increase the push schedule to whatever was appropriate – since it would be automatic. But the temporary situation was a good proving ground for a professionally managed environment.

CMS Requirements

Armed with the detailed knowledge of content generation and processes around the Getty gained during the redesign, the Web group drew up a requirements document for a CMS. Our 'big ticket' items are probably consistent with many museums' CMS wish list, and these were:

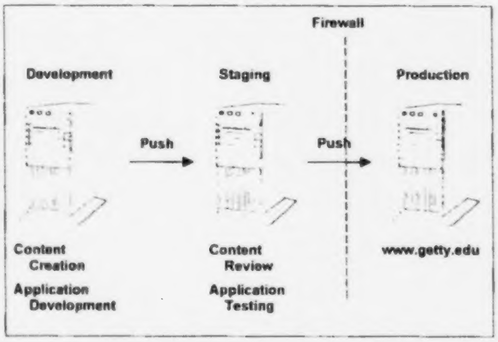

Fig 2: Three-tier Staging Environment

- Template Driven: To guarantee consistent visual design, to separate content from design and to allow multiple content use.
- Used by HTML Illiterate: To allow an interface that lets non-HTML contributors submit and review web content.
- Workflow: The processes established at the Getty to get content to the Web site range from a simple two-step workflow to complex multi-program, concurrent-task processes.
- File Versioning: The ability to track individual file changes, what and by whom.
- Rollback: The ability to publish an edition of the Web site on any previous date.
- Open/Standard Architecture: The Getty's Web infrastructure is based on Unix servers, iPlanet Web servers, Java/perl CGI applications and Oracle databases, all sitting on a Novell network. Any Web application has to live in this environment.
- Virtualisation: The ability to see any changes or new content in the context of the entire Web site – before it is deployed.
- Multiple Web site Management: For getty.edu and for our Intranet GO (Getty Online).
- Cross-platform clients: The Getty has Mac and PC-based contributors.

- Multi-channel Deployment: Eg. XML and WML delivery.
- In-house Maintenance and Development: Not to have to go back to the vendor.
- Vendor Stability: An important consideration given the current economic times.

After tracking the CMS market during the re-design, the Web group focused on three CMS products after the usual RFI (Request For Information) and RFP (Request For Proposal) processes. These were: Stellent's Expedio (http://www.stellent.com), Vignette's StoryServer (http://www.vignette.com) and Interwoven's Teamsite (http://www.interwoven.com). After a bake-off, the Getty selected Interwoven as its preferred CMS vendor.

Implementation

Teamsite is a Web-application suite based on C++, java, perl/cgi and XML technologies. It is a client/server environment with all users interacting with it via a browser. After some testing we established that Internet Explorer was more compatible. The only exception to the browser interface is the module to create workflows – this is a Windows application. Teamsite offers fully integrated Templating, a very flexible and fully integrated Workflow engine, full Virtualisation and integration with a broad range of Web content creation applications such as Dreamweaver and Homesite. It also stands as a solid platform on which to further develop our Web site, by integrating with an array of application, personalization and syndication servers – which are long-term goals for the Web site.

Interwoven's terms and conditions of sale include an agreed method of installation and configuration. This means you need to buy their professional services or contract with a registered Interwoven partner – a so- called 'enabler', of which there are many and for which there is a 'corkage' fee to Interwoven. After reading a number of 'dot com' elegies, we were keen not to have any one of these 'enablers' come in, install the product, then hold us to ransom over professional services to develop our site. Also, our Web group has a strong technical component and should be more than capable of maintaining the environment. The best integration solution for us was Interwoven's so-called 'Fast Forward Core Pack' – a 30-day engagement with a Teamsite consultant and project manager. At the end of this engagement, the client is 'guaranteed' to have a correctly installed and configured product, at least one site deployed through Teamsite as a file-managed process, and up to two templates and two workflows in place.

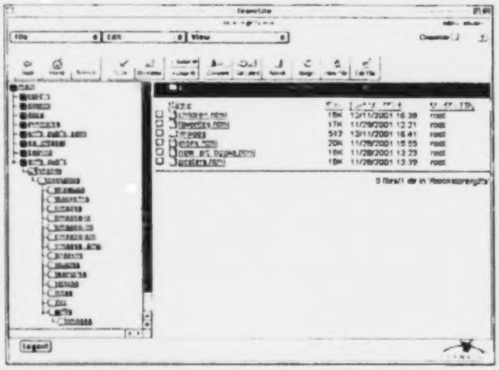

Fig 3: Teamsite's Default View of a Website

We scheduled technical training to coincide with the arrival of the implementation consultant so that we could speak the Teamsite language and could assist in the installation and configuration. We installed the suite onto our Unix staging server, since once operational, we eliminated the need for a separate staging environment.

Environment

Teamsite uses a server's native file structure to manage and deploy content, and manages separate Web sites as branches. Because of this approach, the default view of a Web site within Teamsite is as a file system like Windows Explorer:

Each registered user has an account based on the server's native authorization technology. In our case this is LDAP. Within Teamsite, each user is assigned at least one workarea and has a virtual view of the whole site. This workarea corresponds to the section of the Web site on which they work; they only have permission to modify documents within their own workarea. Every user is also assigned a role depending on contribution level. This determines the authority they have within the environment: author, editor, administrator or master:

BEST COPY AVAILABLE

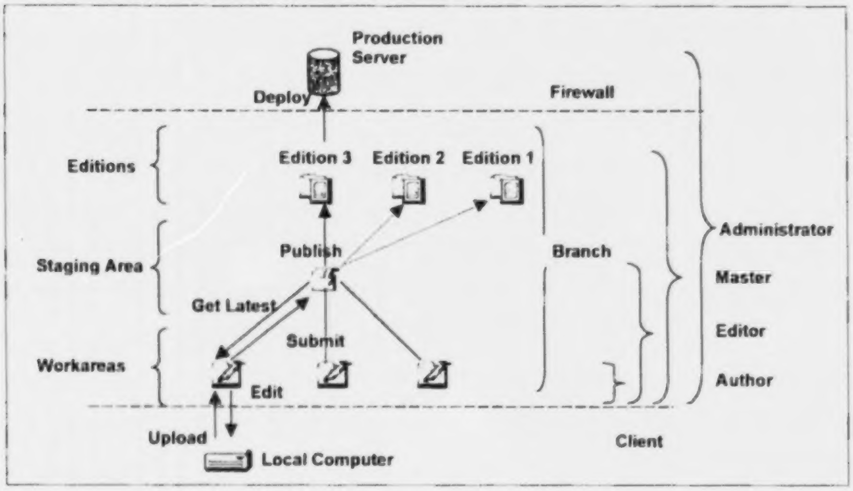

Fig 4: Schematic of the Teamsite Environment and Associated Roles

- Author – the primary content creator: he owns content; can create and edit files; can receive assignments through the workflow engine and has work approved by his workarea owner.

- Editor – creates content and oversees content creation: he owns a workarea; can create, edit and delete files; approves or rejects work of the authors; can submit files to the staging area and can publish editions.

- Administrator – manages an area or branch (Web site), of which he is the owner and can create and delete workareas.

- Master – has absolute power and fundamental administrative control over the product – a role reserved for a Unix System Administrator and a person to be chosen wisely.

Working within the workarea, a user first synchronises with the latest version of content by performing a get latest function within the staging area. When the appropriate edits have been made and checked within the virtualized view of the entire site, the content is submitted back to the staging area. Browsing in the virtualized view of the Web site is no different from browsing the site under normal circumstances. The staging area is where all the content from different workareas is integrated and tested. It is a read-only environment where nothing can be edited. According to the publishing schedule, a snapshot is made of the staging area creating an edition, which can then be deployed to the production server.

We were looking to manage two Web sites with Teamsite: getty.edu and our Intranet, GO. After a week of analysis we began planning two development environments, each of these becoming a branch within Teamsite. We then 'sucked in' each site – a tar'ing, FTP and untar'ing process – which took about a day each. So, within two week' of the consultant's arrival, we had set up a pre-production environment for both sites. With these environments in place, we scheduled some intensive training on the core Teamsite functions of Templating and Workflow – unfortunately, this required a lengthy field trip to Interwoven's training facilities in San Francisco.

Templating

Templating is a core function of a CMS. It allows the separation of content from design, the automatic generation of an HTML page from content, guarantees consistent visual design, and allows Web pages to be generated by HTML-illiterate contributors. A fully templated site is one significant goal of full implementation. Teamsite's templating system is based on XML and requires a thorough understanding to develop templates. We began by analyzing our Web site to identify a list of initial template candidates. We started with the broad assumption

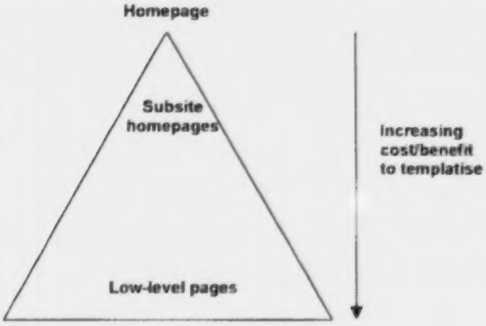

Fig 5: Identifying Template Candidates

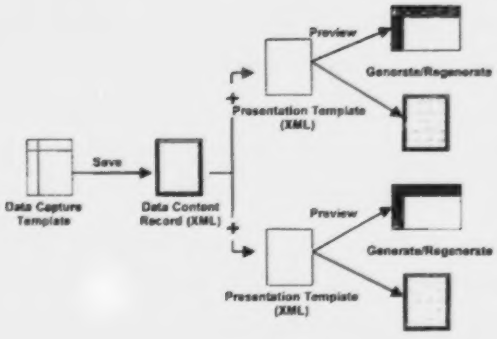
Data Capture Template
Save
Data Content Record (XML)
Presentation Template (XML)
Preview
Generate/Regenerate
Presentation Template (XML)
Preview
Generate/Regenerate

Fig 6: Templating Overview

on the cost/benefit of converting pages under the homepage. This indicated that the greatest cost/ benefit was to be had at the lowest levels, where many pages would be served with a single template. For example, the museum collection subsite consists of approximately 18,000 pages,.Analysis of the design and layout indicated that we could probably serve this with fewer than five templates, possibly even one.

We then applied a number of qualifications: Do a number of people contribute content to this page? Does the content on this page change often? Additionally, we wanted the first template candidates to also establish precedents for resolving issues that would globally impact the site. We focused on two candidates: the Job Postings, which have a lot of similar pages and a daily turnover, and the News Articles, which have a weekly turnover but a significant number of contributors (and consequently a complex associated workflow).

Templating Process

Templating is a three-step process of data capture, storage and presentation, figure 6 shows a schematic of this process.

When a user wants to submit content using a template, they first fill out a data capture template (DCT). This is in an XML configuration file, created by a developer with a good understanding of XML and some basic programming skills. Submitting a completed DCT saves a data content record (DCR) which is either saved as a file or to a database. The last element is a presentation template (PT), another XML file with a variety of markup flavors: ASP, JSP, WML or HTML and also embedded perl callouts and conditional programming tags. The PT is essentially created from an existing HTML page by giving it an XML 'wrapper' and substituting the content areas with appropriate XML tags that are defined in the DCR. Figure 7 shows code fragments tracing a single piece of content, a job title, from capture to presentation:

When the presentation engine combines the DCR and the PT it executes any perl and any conditional tags to generate the final page. The jurisdiction of the different roles in the templating process can be summarized as shown in figure eight.

The templating architecture makes for a very powerful and flexible process and allows for single presentation templates to account for a wide range of related Web pages. Much of the template process can be automated; for example, workflows can be invoked at the DCR commit stage, or the final Web page can be generated automatically when a DCR is saved. It is wise not to underestimate the amount of work required to convert a Web site into templates.

Workflow

Workflow is another core function of a CMS – it defines the content creation and approval processes.

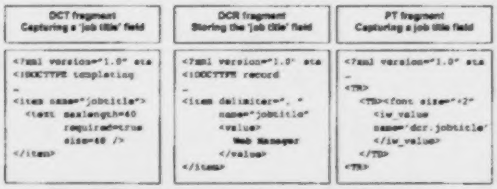

DCT fragment
Capturing a 'job title' field

```
<?xml version="1.0" eta
<!DOCTYPE templating
-
<item name="jobtitle">
  <text maxlength=40
        required=true
        size=40 />
</item>
```

DCR fragment
Storing the 'job title' field

```
<?xml version="1.0" eta
<!DOCTYPE record
-
<item delimiter=", "
        name="jobtitle"
        <value>
          Web Manager
        </value>
</item>
```

PT fragment
Capturing a job title field

```
<?xml version="1.0" eta
-
<TR>
   <TD><font size="+2"
      <iw_value
      name='dcr.jobtitle'
      </iw_value>
   </TD>
<TR>
```

Fig 7: Tracing a single piece of content

BEST COPY AVAILABLE

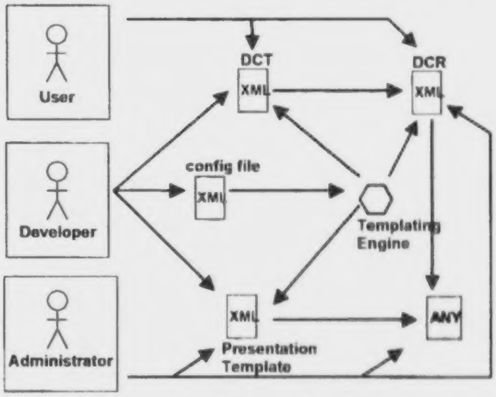

Fig 8: The Jurisdiction of Roles in the Templating Process

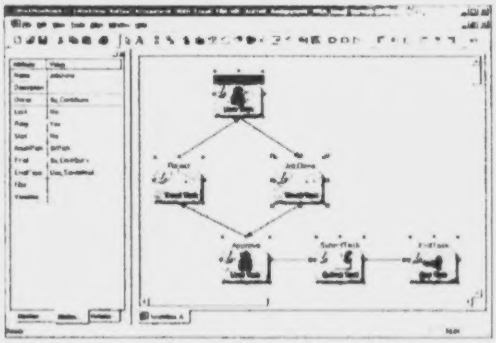

Fig 9: An Example Workflow in WorkflowBuilder

The workflow module in Teamsite was one of the most comprehensive and flexible we reviewed during the selection process. The biggest challenge we found in implementing workflows within the Getty is the analysis of the 'business' processes, which continues to be a challenge. This analysis is a time-consuming practice in an environment where these processes have grown organically over time. When interviewing staff around the Getty, we often found ourselves giving advice on consolidating their processes, since no one had thought to review what was happening. Often when we flowcharted the program's processes, staff were surprised at how redundant their activities were. The concise steps to analyzing a process that we use are:

1. State the general requirement for the process

2. Identify the tasks that need to be performed

3. Identify who needs to perform the tasks

4. Order the tasks in which they need to be performed

5. Review with the process requirement

Having generated and signed off a flow chart representing the workflow process, the next step is to implement the workflow. WorkflowBuilder is a drag-and-drop flowcharting application based heavily on the Visio interface.

Figure 9 shows a typical workflow of six tasks. When invoked within Teamsite, the instance of this workflow is termed a job. It has a variety of variables associated with it, which are defined when the job starts. For example, one variable might be the name of a file that needs updating or the name of a directory where a new page needs to be created. The first task in figure 9 is a content-creation task which will be assigned to a user using the owner variable. When the content has been written and submitted, the workflow sends an e-mail to the appropriate approver with review instructions and a hyperlink to the relevant content. A rejection generates an e-mail back to the owner, but an approval submits the content which might invoke the automatic generation of a Web page ready for deployment. This workflow is very generic. The use of variables allows for many instances of a single workflow – the variables simply travel along with a particular workflow instance. The power of the workflow engine can be extended immensely by additional tasks such as the CGI script invocation and an external task which runs any external application or script on the host machine.

Deployment

Deployment of content is the final step in the content generation process. It is the sending and receiving of files from the base server to the receiver host. Any number of deployment scenarios can be created, based on three basic themes: Deploy files in a specified list – the list can be generated programmatically; Deploy files in a directory by file comparison between the source and the target; Deploy directories by directory comparison between the source and the target. The file transfer

BEST COPY AVAILABLE

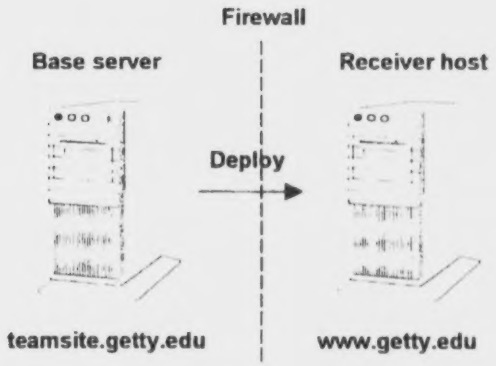

Fig 10: Deployment

executes in a transactional mode, meaning that content will not be replaced on the production server until all the requested files have been received.

The Rollout

One important factor in selecting this product was its ability to manage a Web site at a file level. This allowed us to implement Teamsite without disturbing the contributors around the Getty using their Dreamweaver and Homesite applications. We simply continued to import their files from the development server into Teamsite until we were ready to schedule their training and register them as users. The go-live date was an uneventful day. One day we were 'pushing,' and the next day we were 'deploying'.–This fact alone makes it a resounding success. The process of trawling through the Website leaving templated content in our wake is ongoing, to be completed around June 2002.

Future Initiatives

We selected Teamsite with a long-term web strategy in mind. We chose this particular product because we require a platform that will foster the best opportunities to meet our long term goals and also one that is flexible enough to integrate with presently unknown program initiatives and future www trends. There are two immediate initiatives that are of note:

Vocabulary-Assisted Metatagging – There is a significant metadata initiative at the Getty and extensive research into vocabulary-assisted searching. Currently, searching on getty.edu results in an 'invisible' expansion against our ULAN (United List of Artist Names) vocabulary. By way of an example, if one searches for Carrucci on our Web site, 15 hits are returned to Pontormo but no mention of Carrucci – because Carrucci was the birth name of Pontormo as defined in ULAN. While we will continue to expand search queries against our vocabularies (Art & Architecture Thesaurus and the Thesaurus of Geographic Names are planned shortly), we plan to integrate more fully the metatagging of content at the creation stage. As luck would have it, Teamsite has a module called Metatagger, which we have already begun to investigate.

Personalisation – The 'next big thing' for our Web site is a venture into the world of Personalization. Again, as luck would have it, Interwoven is partnered with a variety of personalization server providers. At the time of writing, we are in the selection process, so stay tuned.

Conclusion

Teamsite is an expensive solution to enterprise-wide content management that the Getty is in the fortunate position to be able to afford. Moreover, our Web strategy is such that we require an 'industrial grade' CMS such as this to achieve our goals. The implementation involved a significant amount of planning, strategizing and work plus a high level of technical skill and support. The ongoing maintenance and development of the site through Teamsite has greatly increased the skill level of the Web group and shifted the typical requirements of a 'manually' maintained Website. In some areas it has polarized the webgroup resources, requiring more technical people to work on the 'back end' but fewer technical and fewer Web-literate people on the 'front end'. The executive-level expectation of a CMS is often a reduction in resources required for Web site development. To a certain extent this is true; however, those reduced resources are at a higher technical level that may offset any perceived budgetary reductions. For us, the webgroup is just as busy after the implementation. The difference is that we're accomplishing more.

Here and There: Managing Multiply-Purposed Digital Assets on the Duyfken Web site

Marjolein Towler, CONSULTAS Pty Ltd., Valerie Hobbs and Diarmuid Pigott, Murdoch University, Australia

Abstract

In this paper we describe a model for managing a dynamic Web site with multiple concurrent versions. The voyage of the Duyfken replica from Australia to the Netherlands has generated great interest in its associated Web site, and it is now planned to extend the original site to mirror sites in both Australia and the Netherlands, each completely bilingual, resulting in four different virtual Web sites. However, the look and feel of the Web site needs to be the same across all the virtual sites, and the media resource used will be the same. We propose a solution based on multiple sets of concurrent metadata, and discuss some general implications of this approach for metadata harvesting and resource discovery. We describe our prototype, which is based on a database that handles parallel sets of metadata, together with the process for accession and annotation of media artefacts, and the dynamic generation of the Web pages from the database.

Keywords: metadata, repurposing, narrative, dynamic Web site, media repository, bilingual

Introduction

The Duyfken Web site is the on-line exhibition of the Duyfken Replica. It serves to tell the story of the conception, construction and embarkation of the replica, and thereafter to recount its voyages around the world, much as the broadsheets informed the citizens of Haarlem and Amsterdam of the progress of the original voyage in the 17th Century. The Web site is the only exhibition available when the ship is between ports, and is currently charting the progress of the VOC2002 Voyagie. It is now planned to extend the original site to mirror sites in both Australia and the Netherlands, each completely bilingual, resulting in four different virtual Web sites.

In this paper we describe the remodeling of the Duyfken Web site to accommodate multiple concurrent versions, of different languages and locations, dynamically delivered to the Web pages from a database and a media repository. The case of the Duyfken Web site provides us with a rare instance where we can examine the theoretical issues involved in multiple concurrency of semantic roles of media artefacts within the context of a practical example with real life pressures and constraints.

We first present a description of the Duyfken Replica project and the parallel development of the Web site. This demonstrates what the situation we are now facing is and how it came to be that way. We then describe the proposed remodelling of the site database and the revised procedures required to create it, drawing on the analysis of patterns in semantic usage and repurposing of media artefacts, and on the ways metadata is used to focus a resource to a goal-driven functionality.

The 1606 Duyfken Replica

In 1606 *Duyfken*, owned by the Dutch East India Company (VOC) and stationed in the East Indies, made a voyage of exploration looking for east and south lands which took it on the first historically recorded voyage to Australia. As part of bringing Australian history to life, the *Duyfken* replica was built at the Lotteries *Duyfken* Village Shipyard in front of the Maritime Museum in Fremantle, Western Australia. Duyfken was built plank first, a method developed by Dutch shipwrights of the 16th and 17th Centuries. No original plans of any ship from the Age of Discovery exist because shipwrights did not use plans drawn on paper or parchment. The only plans were in the master-shipwright's head, and the ships themselves were built by eye. Although replicas or reconstructions of several Age of Discovery ships have been built in recent times, few of them

BEST COPY AVAILABLE

seem to be able to sail anything like as well as the original ships, proving how little we understand of how these ships were built originally.

One of the stated objectives of the *Duyfken* Replica Project has been to produce a reconstruction that sails well enough to emulate the achievements of the original *Duyfken* (see http://www.duyfken.com/replica/experimental.html). So since she was launched, the finished Duyfken replica undertook the re-enactment expedition from Banda, Indonesia, to the Pennefather River in North Queensland, Australia, in 2000 (http://www.duyfken.com/expedition/index.html). She is indeed one of the few square rig replicas of that era capable of making an extended voyage under sail and, according to the known records, lives up to the performance of the original ship. She is currently sailing the original Spice Route in the wake of her 17th Century predecessor on her way to The Netherlands for the 2002 400 year VOC Celebrations (http://www.duyfken.com/voyagie/index.html).

Duyfken Web site - original

The Web site www.duyfken.com, built by Marjolein Towler and her team from Consultas Pty Ltd, is the on-line exhibition of the Duyfken Replica. It was initially developed to set the record straight. Four Duyfken Web sites had sprung up over a period of 6 months after the laying of the keel. None of these were under the control of the Replica Foundation and in some cases were simply factually inaccurate. The Web site was built to make correct information available about the original ship and its historical significance. Collaboration between developer (understanding technology and communication through digital media) and content expert (understanding on a deep level what is relevant to the topic)is crucial to any interactive media (Web) development, and Nick Burningham, maritime historian and the principal researcher and designer of the Duyfken replica, contributed his research and wrote all the original content.

The Web site also kept Web visitors informed of the building progress of the replica. The shipyard was open to visitors, so the physical building process was on exhibit; however, the Web site provided a central point of access to a wider audience, albeit only those with on-line access. The site was also used for Western Australian school curriculum activities, and it built a solid audience of followers from around the world.

The initial Web design was of a simple static HTML page nature. The structure took into account the kind of audiences we thought at the time would be interested in the topic: tall ship aficionados, primary school children, the Dutch and Dutch Australians, fellow maritime archaeologists, the Friends of the Duyfken association, historians, people interested in (model) ship building and ships in general, history, and wood working. This was a very broad audience to please, so although the information structure took the ship as its focal point, we tried to incorporate enough interest for any of the specific audience groups we identified.

The original Web site had two major topics: *Original Duyfken* and *Replica Duyfken*, which each expanded to a number of relevant subtopics. In addition there were general topics available such as Seafarers Links and a Sponsors page. Any visitor arriving at the homepage

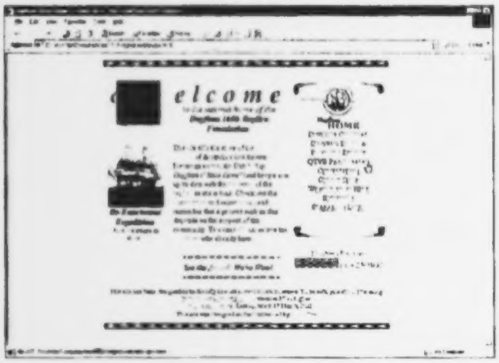

Figure 1 : Original Duyfken Web† site homepage -1997

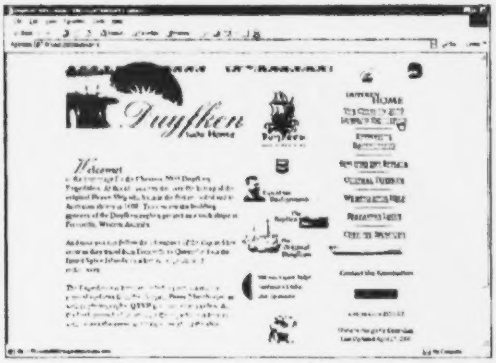

Figure 2 : Duyfken Homepage after Chevron 2000 Expedition was added - 2000

BEST COPY AVAILABLE

was presented with an introductory paragraph and could then traverse to either topic or all subtopics from every page.

To accommodate the different learning styles of users, we incorporated iconographic as well as text links. The screen layout was designed to fit on a 640x480pixel screen display and we limited the colour use to accommodate 256 colour displays. No moving images were added, since from a technology perspective we needed to accommodate low-end receivers and slow modem connections. We even limited the use of frames, because older Web browsers did not support them. The site was designed to look identical on both major browsers of the day (1997): Netscape and Internet Explorer, and we recommended version 3.0 or higher for both.

The only additions that were made to the Web site content in that time were the photographs of the replica building progress and after much agonizing - because it required a plug-in - QuickTime VR panoramas of the same topic. These were produced by Geoff Jagoe and Barb de la Hunty from Mastery Multimedia who offered their expertise. Collectively we argued that being able to see the ship through the panoramas did significantly add to the experience of the Web visitor, enough to justify the wait for the plug-in download, anyway. For those users who did not want to go through with that, there were enough static images readily available.

Content changes and additions had to be made by the Web developers, because it required HTML editing. This was limited to once a month.

duyfken.com Chevron 2000 Re-enactment Expedition

After her launch, the ship took off on the Chevron 2000 Expedition. This re-enactment journey took her from Fremantle, Western Australia, to Jakarta, Indonesia. From there she traveled the original journey to the Pennefather River in North Queensland via the Spice Islands (http://www.duyfken.com/expedition/index.html).

This dramatically changed the requirement for the Web site, because the ship cannot be a physical exhibit when between ports. It thus became clear that the Web site is the *only* exhibit when the ship is at sea, and the only means of communicating her whereabouts to an audience. The kind of audiences interested in such a Web site expanded too. Adventure seekers, armchair travelers, sailing aficionados were added to the list. The site underwent a major overhaul in its layout and interface, partly to accommodate the Expedition section and partly to signify change, so regular visitors would realise that there was additional material.

The primary focus became the Expedition; the original content remained but became ancillary material. The navigation changed to reflect the different choices available and the Expedition section of the site was designed to look significantly different to indicate its separate status. However, we did make the clear choice to maintain the overall Web site homepage, and not make the Expedition homepage the main entry into the site. Our arguments were that the site was still the Duyfken Foundation Web site, not just the Expeditions, and that new visitors would not know of the entire topic choice if all they were presented with was the Expedition homepage.

We also made a concession to technology developments: the screen size grew to 800x600 pixels and the colour settings to 16bit. After 3 years, nearly an eternity in technology terms, we felt that we would not leave too many users behind by doing this.

The content needs changed dramatically too. The project director Graeme Cocks wanted to explore the idea of incorporating the Captains Log on the Web site with the potential to update it daily. This was the basis on which the dynamic database model was brought in. The captain of the ship at the time,

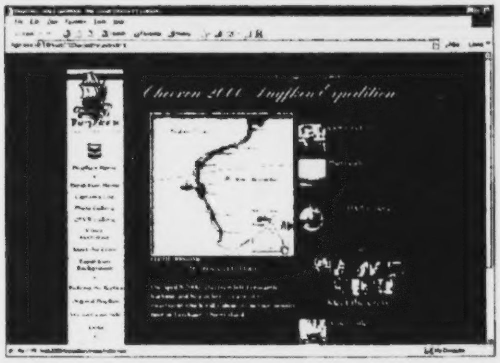

Figure 3 : Chevron 2000 Expedition homepage - 2000

BEST COPY AVAILABLE

Peter Manthorpe, e-mailed his logs via satellite to the land station, the Duyfken Foundation office in Fremantle. Anybody at hand there could and would go simply on-line to a secure URL and upload the logs into the database. Photo images of the journey were sent intermittently either when a reliable phone connection could be made, or when exposed film could be sent to Fremantle for development and scanning. Once digitized the images were uploaded in the database using the same procedure. This also happened to the QuickTime VR panoramas that Nick Burningham produced during the Expedition.

A bulletin board was added to the site to allow site visitors to leave messages. This system was set up to be moderated, but that was never explained well by the programmers, so it did not prove very successful. The project director used it a number of times for uploading press releases.

The database was built to receive the messages, but it was rather unyielding. The more Log messages were loaded, the longer it took to upload. In the end it was so slow the programmers of the system had to come in to change the time out settings on the server. It was not user friendly to ordinary modem use, let alone slow connections in faraway places.

From a design perspective, there were limitations too. The database structure required page templates to be developed. Additional sections were not so easily incorporated, especially not if they required dynamic access, which needed specialist programming. It also required an ISP that allowed client software to be placed on their servers, and that limited the choice of ISPs available.

The advantage was that anybody could upload text and image material through the administration console; it did not required any specialist knowledge of

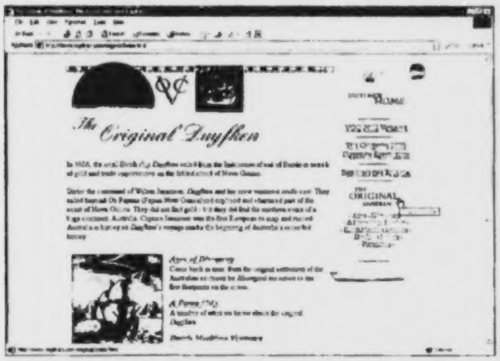

Figure 5 : Homepage for Original Duyfken section 2001

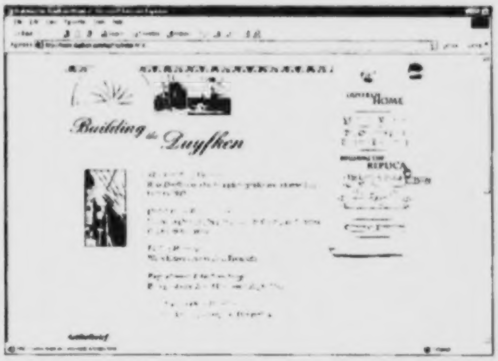

Figure 6 : Homepage for the Building the Replica section - 2001

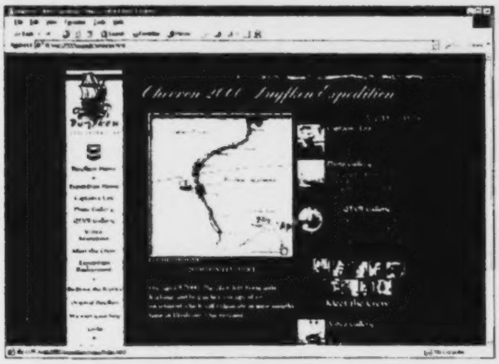

Figure 7 : Homepage 2000 Expedition remained unchanged 2001

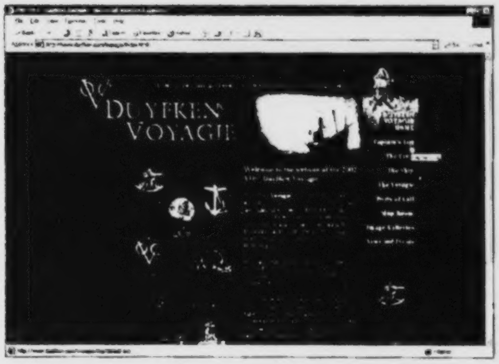

Figure 8 : Homepage VOC2002 Voyagie section - 2001

BEST COPY AVAILABLE

HTML or FTP. This was very important for the Foundation, since as a non-profit organization it was largely dependent on volunteers.

duyfken.com VOC2002 Voyagie

All the dynamic content in the Expedition database was transported to a static HTML format. This made that content more stable (although it increased the size of the Web site to around 600 pages). It also cleared the way to set up the new Voyagie section with a database connected through ASP, which can be edited dynamically through an HTML administration console on a secure URL.

The Voyagie section has five dynamic sections: Captains Log, News & Events, Image Galleries, Whats New and the Dutch version of the Captains Log (Scheepsjournaal). Again, as with the Expedition section, the administrator can upload any data (text, images, QuickTime VR panoramas) through this console to those sections, including the daily logs the current captain, Glenn Williams, emails via satellite.

The database underwent a redesign from the original Expedition one. It was set up to allow for smaller monthly sections to be made. This limits the data entries to no more than 31 (the maximum days in a month), thus limiting the slowdown effect too much data had on the upload procedure. It also not only has a secure URL, but now also requires password access to avoid any accidental or malicious access by unauthorized users.

The interface of the dynamic data pages also changed due to feedback from users. The Captains Log pages of the Expedition were on a dark background, which was very dramatic and aesthetically pleasing as well as easy to read from the screen. However, most visitors wanted to print the logs out for keeping, or reading on paper, and we accommodated that by

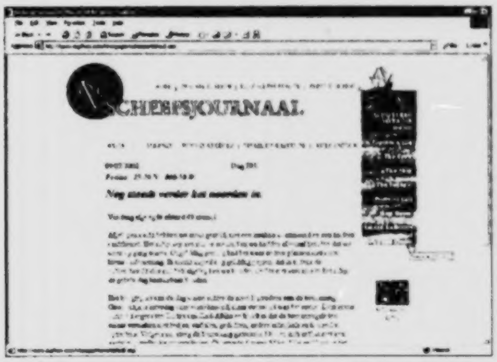

Figure 9 : Dutch translation of Captain's Log - 2001

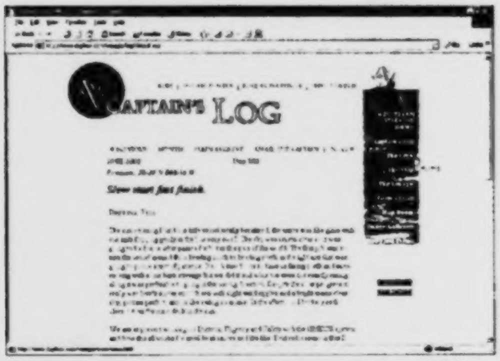

Figure 10 : Captain's Log main page - 2001

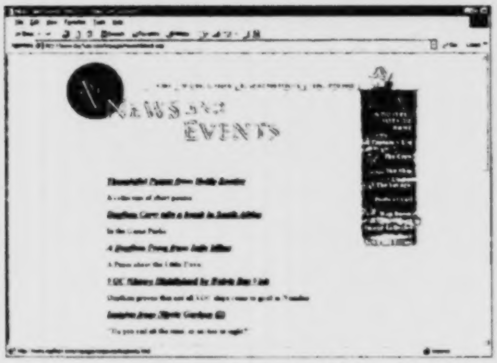

Figure 11 : News & Events main page -2001

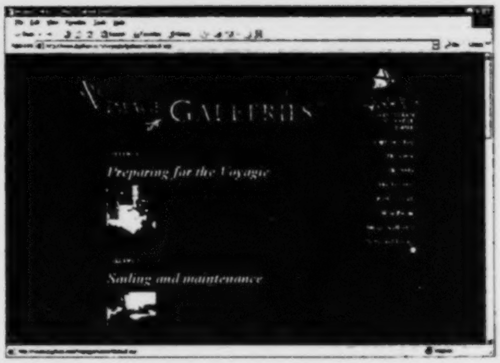

Figure 12 : Image Galleries main page - 2001

BEST COPY AVAILABLE

developing a light background for the Captains Logs in the Voyagie section.

The VOC2002 Voyagie brought an even greater significance than was already there to the Dutch connection, and there was an extensive debate about making the site available in Dutch. At the moment the only Dutch part on the site is the Scheepsjournaal, the Dutch translation of the Captains Log. The Duyfken is sailing to the Netherlands and will be traveling to all old VOC ports when she is there. There are great VOC celebrations planned in the Netherlands throughout the coming year to commemorate the start of the VOC four hundred years ago. All this has posed some problems around control over the database, protocols of access to the database and the way the information is displayed to the individual Web visitor.

The Current Site and its Challenges

The Web site now contains of a mixture of archival, relatively static material retained from the original site, and the extremely dynamic Voyagie section that charts the progress of the replica day-by-day, providing an ongoing narrative of the expedition. Following the great interest generated in the voyage, the Duyfken Foundation wants to translate the entire site into Dutch, and in addition, for speed of access reasons, to create a mirror site of the entire development in the Netherlands.

The site currently consists of a single English language site sitting on an NT server at Web Central, an Australian ISP. The proposed changes effectively mean doubling the site and having both Dutch and English versions sitting side-by-side at Web Central, and a mirror of the site (Dutch/English) on an NT Server at an ISP in the Netherlands. The language issue dictates that there would be more than one person with access and in more than one location (Australia/Netherlands), and there could well be a need to have slightly differing information, especially in News & Events once the Duyfken has arrived in the Netherlands. Thus, instead of a single site, we are looking at having four different virtual sites, each potentially with differing content: English language in Australia, Dutch language in Australia, English language in the Netherlands, and Dutch language in the Netherlands.

All this has serious ramifications for the dynamic uploading of data through the administration console. With the division of the single site into four concurrent virtual sites, there is the potential for problems with control of content and media: it is essential that there be some sort of process control to ensure that one site does not lag behind the others, and that one site cannot be updated without the others, both in terms of copy (textual content) and media content (still images, movies, logs and so on).

This type of problem is normally addressed in software engineering by a version-control system, and for the Duyfken Web site we require a site-versioning system that permits multiple concurrency. What is needed is a system that not only enables the recording of files and their provenance and physical usage but also manages and preserves their contextuality and their semantic usage.

The Problem of Multiply-purposed Media

While the revised Web site has at its core a transfer to a fully data-driven solution, its defining feature is that it is a metadata-driven solution. If a Website is to fit into the general milieu of RDF in the CIMI schema, the appropriate designation of semantic metadata to the media resource is essential (Dunn, 2000 ; Perkins, 2001) . On a static Web site, the metadata is assigned to the entire page, even if the significance of the metadata is the denoting of particular media resources within that page. In the context of dynamic delivery from a database and a media repository, such as we have here, the metadata must be preserved and distributed at artefact level, if individual items are not to be bibliographically lost.

The critical moment in the life of a media artefact is when it is semantically bound, either directly to a purpose or by proxy to a purpose through a document such as a Web page. This is the only point at which intelligence is applied in the form of the skill and judgement of the curator. As media artefacts are at best only minimally self-referential, it is this act of binding that enables automatic sorting, selection and retrieval of sets of artefacts within the repository. The metadata derived from the use then stands for and represents the artefact within the system.

And yet, paradoxically, this act of representation is also necessarily a lens of distortion (Lagoze, 2000) , and properly so - this is how we can focus on salient details at the expense of details of generalities. But this act of distortion means that in repurposing, the specific semantic role metadata cannot be reused, because there is no logical reason for assuming that this beneficial distorting of reality will be equally beneficial to the role anticipated in the repurposing. In other words, although the usages are themselves pertaining to the same artefact, they represent different aspects of reality.

Moreover, every subsequent reuse modifies the net effect of the pre-existing usages. We have argued elsewhere that there must be a set of protocols established in repurposing to ensure that meanings stored in multimedia databases do not drift (Pigott, Hobbs, & Gammack, 2001) .

There are two aspects of multimedia repositories that have implications for metadata practice and artefact repurposing. At system level, we can see that the repository has a telos, or purpose, and intended audience, within its overall theme: adult and child versions of a museum site, or subscriber and free versions of a digital library; or here, the four virtual sites of the Duyfken project. At artefact level, we can consider the usage pattern that the artefact follows throughout its life, through use, repurposing and archiving, which has implications for the extent to which usage is generic or specific, or lying at some stage in between.

These twin dimensions require that we have two aspects to the metadata also: the *teleological* to target the intended audience, and the *abstractive* to ensure the correct level of specificity is used.

When we consider the teleological aspect, we can see that for each instance of metadata usage, there must be a prescribed context of metadata usage, which gives explicit instructions on what constraints are being used, and how these rules are to be applied. This we can term the *metadata frame of reference*, as that is what gives a meaning to the individual usage. A proper program of metadata usage therefore involves the objects for annotation, the instances of annotation, and the rules and keyword sets (the authority set) to use for that annotation.

When we consider the usage pattern aspect, we see that when a media artefact first comes into contact with the repository, it is the decision of the curator that determines its usage pattern, and hence the specific or general nature of the metadata required. There are three main patterns:

- The *immediate permanent use* pattern is where an artefact arrives and is immediately put to use. In this case, it must be catalogued with metadata specific to that use.
- The *immediate archive* pattern is where there is no immediate use for the artefact, but it still has the potential for some later, unspecified usage. Here the metadata needs to be general to enable the artefact to be retrieved for a variety of purposes, none of which is known at the moment of accession.
- The *immediate transitory* use pattern is where the artefact is to be used for some short-term specific purpose and then archived. Here the curator needs to strike a balance between the short-term specific and long-term generic metadata.

There is a special case of usage pattern that we meet in the Duyfken project: the *immediate concurrent use* pattern, where an artefact has more than one semantic purpose *at the moment of accession* (permanent or transitory). These parallel meanings will remain linked through at least the initial part of the artefacts life. Here we need to have multiple sets of specific terms, one for each purpose, and to be able to manage and exploit that relatedness within the overall systematic of organization and retrieval.

When we superimpose the usage pattern of the artefact on top of its teleological aspect, we meet with the problem of continuity and change described in Pigott et al. (2001) , where we discussed the survival of authorial intent; and in Gammack, Pigott, & Hobbs (2001) . If early binding to a specific use requires specific metadata, any later repurposing to another specific use will require metadata specific to that use, a situation which is at least potentially conflicting. If the artefacts have been catalogued with generic metadata, however, a later specific use can be accommodated. This is exactly the problem that we must manage in establishing the systematics of the virtual sites of the Duyfken Web site.

Parallel Metadata: a Prototype for the Revised Duyfken Web site

Our solution is built on a database that lets us represent concurrent use in the Web site, in the form of parallel metadata for media artefacts stored and maintained for the four virtual sites, so that the correct version is available to be called upon for a particular site and language. Over and above this base requirement, we need to ensure that we represent the components in such a way that the site of which they are a part can preserve its narrative structure, in order that we may call on those resources in the correct context.

We therefore have several clear objectives for the prototype:

1. It must retain the existing style of the site, and permit it to expand to four separate, yet parallel, site.

2. To accomplish this, it must permit separate sets of metadata for each of the four concurrent usages (English languageAustralian site, EnglishNL, DutchAustralia, DutchNL).

3. This metadata should fit into the industry standard practices of defining and exploiting metadata in order that the local expenditure of effort have a universal return.

4. There should be a naturalistic process for the accession/cataloguing of media, to permit an unobtrusive capturing of the metadata and its storage for the appropriate virtual site.

5. There should be a process for generating the Web pages in context that gives rise to the dynamic delivery of the four separate virtual sites, a process which should also ensure the narrative of the particular site is preserved.

When we modeled the database, we identified four main entities of interest: Event, Setting, Place and Agent. The data model centres on the concept of an EVENT. Events are things that happen; on the voyage, or before or after it stages in building the replica, the departure and arrival of the ship, individual passages and stays at ports, and the news and events, including captains log, that are mailed regularly to the Web site. Events are linked with other events through a recursive relationship, as an event such as a stay at a port is also part of the larger event of the relevant passage, which itself is part of the entire voyage. We also note that since the Web site is itself part of the events of the voyage, charting as it does its chronology and narrrative, the construction of each Web page will also be recorded in the database as an event, and the Web pages themselves stored as artefacts.

Events occur at Places, such as ports or unnamed locations of latitude/longitude; and in a particular Setting, which could be on the ship, a particular part of the ship, or onshore. Events have Agents associated with them. Agents can be individual people, crew positions, or collective groups such as crew or land staff. Again, each of these entities is in a recursive relationship with itself, as each place or setting may be a part of a larger whole, and individual agents may be part of a group.

The relationships identified among the entities enable the narrative structure of the Voyagie site to be represented, as well as the more static information present in the original site. Since all the main entities are in recursive relationships, the inherent hierarchies in the Duyfken voyage can be modeled. This allows us to extract more complex combinations of information, for example getting connections between legs of the journey and members of crew, via ports of call, and the watches of the night.

In the Duyfken Web site, we have four virtual sites, consisting of the combinations of two languages and two countries. So, for each physical media artefact playing its particular role in the narrative, we need to consider four sets of metadata. We can identify several possibilities resulting from this:

1. The language is different, but the site makes no difference to the metadata that is used. This occurs with the technical metadata, where there is no interpretation involved, and also with the conventional data in the Event, Place, Setting and Agent tables. This situation can be accommodated by having dual language fields within the same record. (Although this is a hard coding solution it is simplest to manage and is justifiable here as it is unlikely that any other languages will be required for the site.)

2. The site at which the artefact is used may impose a different interpretation on it. For example, the same image could be captioned differently in the Netherlands and in Australia. This involves not just language translation, but a different set of semantic metadata. For this situation, we use four parallel metadata records.

3. If the interpretation or translation required for a site involves editing or copying the artefact, we end up with multiple physical artefacts: for example, each entry in the Captains Log/Scheepsjournal is stored as two separate documents. Each document, being a separate artefact, would then require dual language fields for the technical metadata and four parallel records for the semantic metadata, as described above.

Process for Accessioning

The Web site, and therefore the database which underpins it, is a live chronicle of events. The selection of media for a section of the narrative is done by a series of page wizards, permitting searching for material either generally or specifically, and from the perspective of an entity or from a unified set based on the metadata. An interesting challenge here is to present multiple possibilities for keyword or character/place referencing without losing the flow of the selection process.

The Web pages are designed to have zones of placement, as close to the current Duyfken site as possible. Use of a cascading style sheet in conjunction with the stored text permits a set of different semantically purposed page regions, each with content that is apposite and current. The use of summary textual metadata permits headlines to be created, and serve as keyholes to the pages where the information is given in full, and also from each section to a set of previous entries for that section. This fits in nicely with the image of the site as a narrative chronology, similar to that in old Amsterdam. There is a strong publishing feel to the site, which enhances its feeling of vitality and immediacy.

The new Web site will have curators in each country whose responsibility it is to add the media and data to the database. Although there need to be four virtual curators, one for each site, the roles could be filled by anything from one to four or more individuals; thus in the following discussion, curator refers to the appropriate virtual curator. The secure access of the existing site will be continued, and curators will be required to log in with passwords.

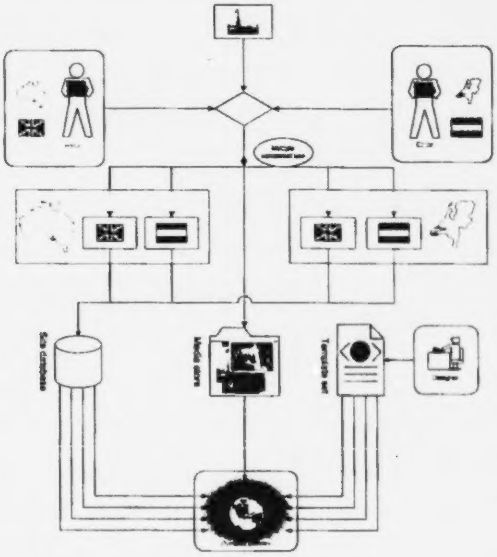

Figure 13: Process for accessioning

Media artefacts and notification of events will arrive regularly, as the voyage progresses, and the database, media repository and Web pages will be updated at the same time, in the context of the ongoing narrative. Figure 13 illustrates the logical process of accession of media and data and the delivery of Web pages.

The new site is event driven, and there are three different kinds of stimulus from the world that the curator must respond to:

- The arrival of new media artefacts, such as images or an e-mail of the Captains Log. These artefacts will always be part of the context of an event, whether new or previous.
- An event happening, such as arrival at a port (which may or may not have accompanying media).
- A general need for updating the site.

Each of these stimuli is a trigger that causes the curator to be notified, using internal messaging, that the system requires attention. We will now describe each of these processes in turn.

When new media artefacts arrive, there is likely to be a surplus of artefacts, and the first judgement of the curator will be one of appropriateness of the artefact (be it still image, recording, e-mail, movie) to the immediate site narrative. In some instances (the Captains Log is the main one) the artefact is immediately needed for all locations (and in the case of the log, it must be translated and a separate document created). With a set of volunteer images of a stay at a port, on the other hand, it may be that only three images out of 20 would be used.

The next question arises as to the usage status and the metadata requirements of the remainder. As discussed previously, there will be a tradeoff between very specific metadata for immediate use, and more general metadata for archival use. The artefacts that are immediately archived are often going to be very similar to the ones selected for immediate use, and a level of description and keywording will be aimed for to permit a general location, but perhaps not the detailed analysis of action and character of an artefact designated for immediate use.

It is essential that every artefact receive its full complement of parallel metadata, and that this process be managed. A system of locking will ensure that it is impossible for more than one curator simultaneously to access and catalogue the same artefact, and a set of pages inserted after the initial accession pages present in the original facilitates the addition of the metadata. The metadata will be drawn from best practice technical metadata requirement and the selected semantic role metadata sets.

The first curator to respond to the message would be the one to add metadata from an agreed set, with place-holding metadata put in the system for subsequent curators. The place-holding metadata will be generic rather than specific, and will be in the appropriate language for the usage, drawn from key-paired thesaurus and similar wordsets, wherein every entry in the thesaurus will have a Dutch as well as an English term. The titles, descriptions and other free text will have to be left with a translator, or team of translators, who will be delegated the individual tasks automatically by internal messaging. The next point in the chain is a form of copying-on, whereby the archived (not immediately used) artefacts are given some of the used, generic, metadata automatically (Figure 14).

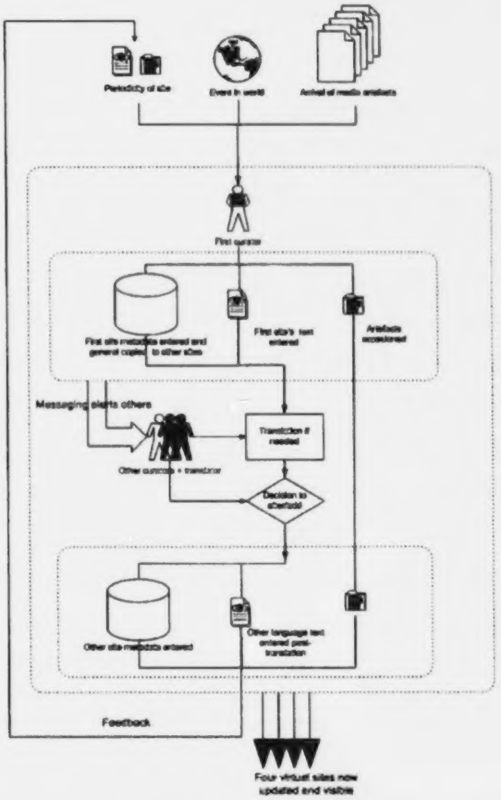
Periodicity of site
Event in world
Arrival of media artefacts
First curator
First site metadata entered and general copied to other sites
First site's text entered
Artefacts accessioned
Messaging alerts others
Other curators + translator
Translation if needed
Decision to default
Other site metadata entered
Other language text entered post-translation
Feedback
Four virtual sites now updated and visible

Figure 14. Curatorial workflow.

The second curator on the scene follows the same process of selecting from the artefact set and accepting or amending the default metadata, and if necessary adding specific metadata. It may also be that the second curator finds an alternative artefact in the batch to suit their purposes better, in which case, the two sites will not have identical media artefact usage.

The artefacts are all placed into a media repository, which is automatically mirrored, and updated daily. This ensures that there are no bandwidth constraints in the delivery of media for analysis.

As the media artefacts will always added be in a context the context of an event, agent, place or setting the possibility of adding a new record to the database tables at the same time has to be taken into account. Again, the first curator would enter the information, and the other language version would be delegated to the translators.

The second case that the curator must respond to is when events occur and trigger an update to the database but without accompanying media. In this case, the curator may need to search through the existing media repository to locate suitable media artefacts for illustration.

Finally, there is always the need for the Web site to be fresh and interesting for visitors, so in the absence of any new media or events to be recorded, it may be necessary for the curator to write copy on more general themes. This could be prompted on a time or calendar basis. Here the curator may need to search the database for suitable media and perhaps also for stories from the records of events.

After the completion of each update of a Web page at a virtual site, the creation of the page is itself recorded as an event and the Web page is stored as part of the media repository. An interesting side-effect of the Web site creating dynamic pages as a chronicle is that the contents, while dynamic, are representative of their period of currency, and so themselves become the subject of harvesting and storage. It is interesting to consider whether the individual pages will require the direct attention of the editor for adding metadata presumably the requirement for internal monitoring of the site could well have this as a clear benefit, with the response of the site to certain events being indicated with candour for internal purposes where they would not have been evident from the site itself.

What distinguishes the Duyfken project is the need for the parallel versions to remain concurrent in time throughout the constantly dynamic process of updating the narrative. Ideally, updates to the four virtual sites would be handled as a single transaction, so that no version lags behind the others. In practice, however, it may be necessary to set a timeout period so that if for any reason a curator fails to respond to a message by posting an update, the remaining sites are not held up indefinitely.

Implementation

The Web site is currently on a Windows NT server using ASP, with a perl/mysql backend delivering media as required. There is a hybrid ECMAScript (JavaScript) and ASP development environment, showing the different stages of the project development, as described in the introduction. The initial prototype, replacing the Voyagie/ship's journal section, was done with the Pasigraphy scripting environment (developed by D. Pigott) and had the data in Microsoft Access tables under IIS5.

The next stage, the pilot project, makes use of a straightforward ASP/ADO format, designed and managed through Microsoft FrontPage. The interoperability of IIS 5 with FrontPage and the relative ease of use for the non-technical user will make it possible for the curators to experiment with alternative features in relative safety, due to the componentisation of the design. The choice of this platform was made partly for reasons of continuity, and partly with a view to making a generalisable solution that would be usable elsewhere. The tables are currently in a SQLServer7 backend, but upgrading to SQLServer 2000 presents no challenges at all.

The movement of files is done via HTTP Upload to avoid any potential confusion associated with standard UNIX ftp. The exchange of metadata and the new entities is done automatically by time-triggered stored procedures, while the inter-curatorial communication is done by an internal, project-style, messaging system. It is hoped that by designing the system to be scalable, there will be no barrier to repurposing the system itself.

The design of the database was done using the entity-media modeling methodology (Hobbs & Pigott, 2002) and the process of designing for concurrency is described in Hobbs, Pigott & Towler (in prep). An unusual feature of the design is the recursive nature of the four principal entities involved, which has led to an unusually adaptable system for querying and producing network reports. The parallel metadata as required for the virtual sites is stored using dual language fields for technical metadata and conventional data, and separate records for semantic metadata. Keyword sets with linked terms are used to accommodate the complex nature of bilingual synonymy.

In terms of the choice of media artefact formats, although the Duyfken project is to a large degree archival, there has had to be a trade-off between the needs of archiving material (long term, guaranteed readability, multiple store, etc) and the needs of the site, which is to be maximally viewable, with an optimum download. Certainly the Duyfken voyage is of interest to places where it calls on its way, and there is a responsibility of the site to permit easy viewing for those on the other side of the digital divide. Moreover, there has been no way of guaranteeing a single standard for quality where some of the digital material is volunteered, and already arrives in a digital format. The choice to date has been to opt for high quality JPEG formats, QuickTimeVR and HTML4 compliant text. With the transition to the new site, the needs of longevity are probably better served by continuity, due to the massive task of retroconversion. Certainly, embracing formats such as fractal compression or wavelet compression ahead of their general availability in browsers would go against the Web sites principle of universality of access.

The Duyfken Web Site and Resource Discovery

The late-binding use of media and text resources in the revised Web site not only makes for timeliness and accuracy, but also makes it possible to have an inventory of the entire media resource pool, and the benefits that entails. Not only are there savings in time and money through the reuse of media artefacts, but we are also moving towards an environment of repurposing and media discovery, and the increased semantic richness that arrives with that discovery. It is the intention that (though the pages are dynamically created) the Duyfken Web site be a source for Dublin Core compliant harvesting, to permit discovery by other museum and historical presences on the Internet. But in facing up to the challenges this presents, some problems are immediately apparent.

The IFLA FRBR common logical method for metadata presents the information resource in one of four states: Work, Expression, Manifestation and Item (Bearman, Miller, Rust, Trant, & Weibel, 1999) . It is not straightforward to fit the work of the Duyfken Web site into this paradigm, however, for several reasons.

The fundamental nature of the DC paradigm is of works of cultural significance being described and accessed, reflecting its bookish and artistic heritage. This is not readily applicable to our situation. The artefacts are not necessarily of worth in themselves, and there is a documentary rather than an artistic presence on the Web site (this is not to deny the aesthetic worth, but rather to describe its direction). There is no real creator or authorship of importance, but rather the *story* itself is the main bearer. The sailors themselves are telling the story through their actions, and that is the authorial component that should be stressed.

Neither is there any manifestation of importance. The presence of the media artefacts here are all of an illustrative nature: here we are concerned with trying to represent events, and artefacts are chosen for their ability to represent that narrative, rather than for any significance in themselves.

That said, what can we gain from the insight of the OAG and the DC? Certainly the paradigm of interoperability is crucial here: permitting the discovery of information through judicious use of RDF, and through the appropriate access to properly accessioned media artefacts in a repository, is the responsibility of the conscientious cultural store point-of-presence so what can be asked of the site itself in this paradigm? Here we find resource discovery not so much between the harvester and the site components, but as part of the harvesting of the site itself as a resource; so metadata practice should have as its first priority harvesting of the site as a narrative rather than individual points of access to the site.

Perhaps it would be more appropriate to consider the types of requests that might be made of the repository in actual use. Here questions would be asked of the site as of an expert on Dutch maritime history, so the resources could be exposed and harvested in the form of specific rather than general queries. Not so much what media have you available as what can you tell me about this subject in particular? This goes back to our earlier discussion of how metadata information should be prepared for early and late binding use.

Editor's Note: A chronology of the Dufken Web Site can be found in the electronic version of this paper

References

Bearman, D., Miller, E., Rust, G., Trant, J., & Weibel, S. (1999). A Common Model to Support Interoperable Metadata: Progress report on reconciling metadata requirements from the Dublin Core and INDECS/DOI Communities. *D-Lib Magazine, 5*(1).

Dunn, H. (2000). Collection Level Description - the Museum Perspective. *D-Lib Magazine,* 6(9).

Gammack, J. G., Pigott, D. J., & Hobbs, V. J. (2001). *Context as history: the cat's cradle network.* Paper presented at the Australasian Conference on Knowledge Management and Intelligent Decision Support - ACKMIDS '2001.

Hobbs, V. J., & Pigott, D. J. (2002). Entity-media modelling: conceptual modelling for multimedia database design. In D. R. G. Harindranath, John A.A. Sillince, Wita Wojtkowski, W. Gregory Wojtkowski, Stanislaw Wrycza and Joze Zupancic (Ed.), *New Perspectives on Information Systems Development: Theory, Methods and Practice*: Kluwer Academic, New York, USA.

Lagoze, C. (2000). Metadata Musings, *Talk presented at University of Virginia, April 2000.*

Perkins, J. (2001). *A new way of making cultural information resources visible on the Web: Museums and the Open Archives Initiative.* Paper presented at the Museums and the Web 2001, Seattle.

Pigott, D. J., Hobbs, V. J., & Gammack, J. G. (2001). An approach to managing repurposing of digitised knowledge assets. *Australian Journal of Information Systems* (Special Issue on Knowledge Management).

Pyramid Power: A Train-the-Trainer Model to Increase Teacher Usage of the ArtsConnectEd On-line Resource

Scott Sayre, The Minneapolis Institute of Arts, and Kris Wetterlund, Independent Educational Writer, ArtsConnectEd Lead Trainer

http://www.artsconnected.org

Abstract

In 2000, two Minnesota art museums began the development of statewide networks for training teachers to integrate internet-based educational tools and resources into their classrooms and teaching techniques. This paper examines the design and implementation of a train-the-trainer program designed to promote the use of www.artsconnected.org (a Web site developed jointly by The Minneapolis Institute of Arts and Walker Art Center) in classrooms across the state of Minnesota. The goals of the program: 1) create a collaborative laboratory for exploring meaningful classroom applications of on-line teacher resources; 2) create a community that will sustain the use of on-line teacher resources throughout Minnesota; 3) build bridges between the cultures of the classroom and the art museum; 4) increase the sophistication of teachers' use of on-line teacher resources; and 5) expand teachers' use of technology in general. Personal anecdotes and valuable lessons learned through formal evaluation of this program illuminate global issues of interest to all museum educators.

Keywords: teachers, teacher training, K-12 teachers, K-12 classrooms, education, K-12 education, art education, museum education, museum collections, museum Web site, lesson plans, Internet, Web site, partnership.

Introduction

Two Minnesota art museums have developed a statewide network for training teachers to integrate ArtsConnectEd, an internet-based educational resource, into traditional classrooms and teaching techniques. A train-the-trainer program was designed to promote the use of www.artsconnected.org (a Web site developed jointly by The Minneapolis Institute of Arts and Walker Art Center) in classrooms across the state of Minnesota.

The Evolving Partnership

The partnership between the Education Departments of The Minneapolis Institute of Arts (MIA) and the Walker Art Center began in 1995 when each museum designed catalogues of educational resources for teachers and mailed them together in one envelope. The joint mailing required that the museums collaborate on design and pool their mailing lists and share mailing costs. The response from the K-12 teaching community in Minnesota was very positive. Teachers expressed the sentiment that finally the two largest art museums in the state had quit competing with each other and focused on serving teachers.

As the Web emerged, both museums also made their catalogues available electronically through their respective Web sites. Simultaneously, some of the material listed in the catalogues was being converted for on-line delivery. Originally the on-line catalogues and other electronic teaching resources were straight HTML. As interest in these resources grew, so did the desire to make the material searchable. In addition, both museums began to digitize their art collections to participate as founding members of the Art Museum Image Consortium (http://www.amico.org). In 1997, the Minnesota Office of Technology began offering funding to encourage the development of on-line resources. During this year, technology staff from the MIA and Walker met to explore on-line project ideas that might overlap, address State funding requirements, and further bolster each museum's commitment to on-line resources. The Integrated Art Information Access Program, later renamed ArtsConnectEd, was the product of these meetings.

Technical Grassroots

ArtsConnectEd grew to become a portal to the combined digital resources of the MIA and Walker,

Museums and the Web 2002: Proceedings
Archives & Museum Informatics, 2002, p. 137

BEST COPY AVAILABLE

including works of art, educational resources, audio, video and text archives and library catalogues. Funding opportunities as well as the educational resource catalogues that were the birth of the partnership motivated the technologists to consider K-12 educators a primary audience. In 1997, the Minnesota Department of Children, Families and Learning awarded ArtsConnectEd $1 million to continue its work putting the collections of both museums online.

Both the Walker and the MIA were struggling to address growing demands to put traditional print-based educational resources into the hands of educators, and ArtsConnectEd was designed to provide an unlimited 24/7 means of addressing those demands. The museums quickly realized that they would have to make an investment in helping teachers use technology and the Internet. Wiring of schools in Minnesota and nationwide was moving quickly, and ArtsConnectEd developers assumed that access would increase as quickly as new resources were generated. ArtsConnectEd was able to make a case as an investment in technology and education, and in 1999 received an additional grant of $2 million to continue its work for and with teachers.

Interface Design

ArtsConnectEd began with the idea that using query, similar to Google as the primary interface for searching the collections of both museums, would best serve the audience of K-12 teachers. But following usabilty testing, technologists began prototyping new interfaces designed to better suit teachers' needs. They enlisted the help of usability experts and the museum educators at both institutions. The resulting interface, publicly released in the fall of 1999, incorporated resource-specific queries, wizard/assistants and menus to better meet the needs and requirements of the K-12 educational audience.

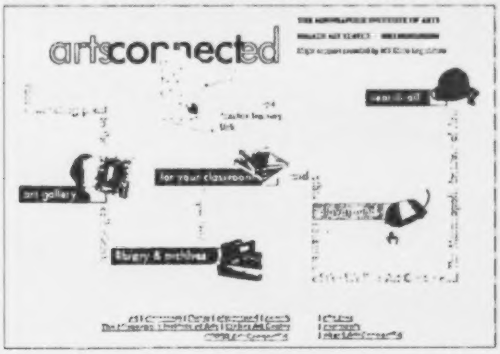

Fig. 1: ArtsConnectEd Home page

Management by Committee

As the usage, care and feeding of ArtsConnectEd began to grow, a more formal team structure was needed to set and obtain long-term project goals. Four committees composed of representatives from both the MIA and the Walker were established to achieve these ends: Steering Committee, Technology Committee, Education Committee and Marketing Committee. The Education Committee defined new content and resources to be made available on ArtsConnectEd and began to explore strategies of training teachers to use it.

A State-wide Training Strategy

Equitable distribution

The state of Minnesota is 84,068 square miles and contains 341 school districts to serve a population of 4,610,000. The MIA and Walker are located in the most densely populated urban area of Minneapolis, but the majority of Minnesota educators outside the Twin City area were grossly underserved. The state legislature wanted their investment in ArtsConnectEd to reach the whole state. In order to secure additional funding from the state, the museums needed to design their teacher training program to impact all of the state of Minnesota.

Part Marketing, Part Training

The ArtsConnectEd Education Committee identified two initial goals: to promote the use of ArtsConnectEd in classrooms around Minnesota, and to coach and support teachers learning to use ArtsConnectEd. Providing teachers time to learn to use ArtsConnectEd in a workshop setting might help address the problems teachers have finding time on their own for such professional development. ArtsConnectEd teacher training could also supply the support that teachers need in their use of technology and relieve school districts of some staff development pressures felt around technology in the classroom.

BEST COPY AVAILABLE

The Train-the-Trainer Model

The ArtsConnectEd Education Committee began its design of a state-wide teacher training program by considering a traveling resource model such as Plains Rolling Art Gallery (http://www.plainsart.org/education/rpag.shtml) or Experience Music's Electric Bus (http://www.emplive.com/visit/electricbus/index.asp). However, this strategy had a "here and gone" impact without lasting support, and one trainer could reach only a small number of teachers. The alternative approach of a train-the-trainer model was conceived, using a pyramid system to reach as many teachers as possible throughout Minnesota.

The train-the-trainer model is based on one lead trainer developing a highly skilled core group of teachers from around the state, who would in turn conduct workshops in their regions. After consulting lobbyists and administrators of both museums, the Education Committee chose to recruit a core group of 24 teachers, 18 from around the state and six from the Minneapolis/St. Paul metro area. Each member of this core group would contract to teach at least four workshops of 15 teachers each after they were trained, so 24 trainers would workshop 60 teachers each, resulting in a grand total of 1440 Minnesota teachers trained to use ArtsConnectEd in their classrooms during the 2001-2002 school year. Theorizing that each teacher is responsible for at least 30 students during a school year, the Education Committee estimated that students using ArtsConnectEd could reach 43,200 during 2001-2002.

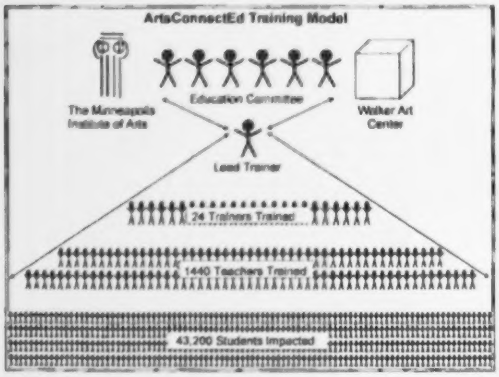

Fig. 2: Training trainers

Project Budget

ArtsConnectEd money was dedicated to three primary areas: 1) the continued digitization and conversion of works of art and archives, 2) the development of new content in the form of on-line educational resources, and 3) the statewide teacher-training program.

While the Education Committee conceived the strategies of the overall training program, its design was facilitated through external contractors. Staffing in the figures below is non-museum staffing. Materials for training are included in the cost per trainer.

YEAR ONE	YEAR TWO
$5375 per trainer just for training	$1695 per trainer for workshops
$45,100 staffing cost	$57,400 staffing costs
$50,000 marketing	

Table 1: Project Budget

ArtsConnectEd Training Plan

Once the model for the training program was determined, the Education Committee examined its resources and decided that each trainer would receive a Macintosh G3 Powerbook computer in order to ensure equal access to standardized equipment and software. Each trainer would also receive reimbursement for lodging and mileage and a $200 stipend to travel to Minneapolis for three training sessions in 2001. In return, the trainers would commit to use ArtsConnectEd with their students in the classroom, as well as plan and deliver four ArtsConnectEd workshops in their regions during the 2001-2002 school year. Trainers would receive a $200 stipend for each workshop they conducted and an overall expense account of $500. During the two-year training period the trainers would also participate in an on-line threaded discussion list with the lead trainer and each other, as well as participate in an evaluation of the program.

Lead Trainer

In the summer of 2000, the Education Committee began the search for an ArtsConnectEd Lead Trainer. The Lead Trainer would be responsible for developing both print and electronic training materials and classroom assessment models, and would moderate the threaded discussion list with the 24 trainers using *WebBoard* software. Once the 24 trainers were fully trained, the Lead Trainer would observe the workshops they provided for their peers and work with a contracted evaluator to study the effectiveness of the training program.

The Education Committee identified the following qualifications for the Lead Trainer: 1) a licensed art educator experienced in teaching with technology and in developing online teacher resources; 2) previous experience working with art museums; 3) knowledge of the Minnesota community of educators; and 4) willing to travel to various parts of Minnesota to observe each of the 24 trainers in action.

Because the Lead Trainer was a temporary contract position, the Lead Trainer's relationship to permanent museum staff would need to be transitional. Therefore, each museum selected a member of its education staff to attend all training workshops and work closely with the Lead Trainer to ensure continuity. The Lead Trainer was given a museum e-mail address (acetrainer@artsmia.org) and an assistant trainer was identified on staff at the MIA to serve as a liaison between the Lead Trainer and in-house museum services and procedures.

Trainer recruitment strategy and application process

The Committee decided to target art teachers in their recruitment strategies in order to promote the art education profession in Minnesota. The Education Committee reasoned that recruiting art specialists to become ArtsConnectEd experts would be a statement of support for the profession and a leg up for art specialists who were often shut out of access to technology in favor of more "technical" subject areas like business or science.

Another consideration was whether to target teachers who needed education in technology and inspiration to use technology in the art classroom, or to target those teachers who were already technology savvy and had demonstrated leadership in the field. In the end the Committee chose to recruit leaders in the field, hoping to inspire others through the examples set by the trainers.

To avoid waste or duplication, a two-step application process was developed. An applicant formally requested an application package, and if still interested, completed the application. An advertisement was placed in folders handed out at the fall conference for Art Educators of Minnesota, and also was sent to the state's mailing list of art educators and the state's Best Practices Network, a group of 50 arts teachers assembled from all over the state who are knowledgeable about the research on best practices for effective teaching and learning in the arts.

The application packet consisted of a letter detailing the commitment a trainer would need to make if selected, and the application procedure. In order to ensure that applicants had the blessing of their school administrator, a letter of support was also required with the completed application. The application form required applicants to specify which grades they taught. The Education Committee had determined that trainers would be chosen from those teaching grades 4-12 since ArtsConnectEd wasn't designed for use by preschool children or the very young grades K-3. Thus an art teacher responsible for grades K-6 in a large elementary school fit the target applicant profile more than a teacher responsible for grades K-4 in a smaller school.

A question was included on the application regarding the teacher's prior experience with museum educational materials. The applicant's demonstrated leadership history was required in terms of memberships in professional organizations and awards or honors, as well as a short essay outlining their philosophy in regard to the integration of technology in classrooms. A series of questions was included about how often the applicant used the Internet to determine the applicant's comfort and familiarity with Internet resources. Finally, each applicant was also required to sign a statement saying that they read and understood the job description, to avoid having applicants who later withdrew over misunderstandings regarding what was required.

Trainer selection

150 requests for applications were received by the December 1, 2000 deadline, and 89 completed applications were received by the December 15, 2000 deadline. Each application went through an initial screening process in which applicants who didn't fit the recruitment criteria of teaching in grades 4-12, and applicants who indicated that they never used museum resources or never used the Internet were eliminated. The geographic "home" of the remain-

ing applicants was plotted on a map of Minnesota. The Education Committee met and reviewed the remaining applications, choosing 24 based on the criteria of geographic location, qualities of leadership, experience training teachers, and on the essays regarding the use of technology in the classroom. As a final check, at least one of the two references listed on each of the 24 finalists' applications was contacted. The applicants who were not selected were sent "thank you" letters and a gift of catalogues of the MIA's collection and the Walker Sculpture Garden.

The selected group represented a distribution of elementary, middle and high school art teachers. All were informed at their first training meeting in February 2001 that they would be required to sign a contract with ArtsConnectEd for their work over two years (see Appendix A).

The Marketing Plan

A marketing firm was hired to develop a marketing plan and materials for ArtsConnectEd teacher-training during year two when the trainers were in the field teaching their workshops. An ArtsConnectEd poster and three-ring-binder were developed before the 2001-2002 school year began, enabling trainers to use posters to market workshops and binders to be filled with customized material for each workshop. A media kit was developed and the marketing firm tracked the workshops as they were scheduled on the threaded discussion list. As each workshop was scheduled, the marketing firm sent a media kit to the host community's press.

Program Goals

Once the 24 trainers were selected the Education Committee interviewed and contracted an evaluator specializing in outcome-based evaluation to design an evaluation of the train-the-trainer program as well as one for the workshops the trainers would teach in year two. Goals established were 1) create a collaborative laboratory for exploring meaningful classroom applications of on-line teacher resources; 2) create a community that will sustain the use of on-line teacher resources throughout Minnesota; 3) build bridges between the cultures of the classroom and the art museum; 4) increase the sophistication of teachers' use of on-line teacher resources; and 5) expand teachers' use of technology in general.

Training the Trainers

Session One

Goals for Session One

1. Increase trainers' computer competence, particularly with their Macintosh Powerbooks
2. Effectively communicate the basic structure and uses of ArtsConnectEd.
3. Build a learning community between the two subgroups of twelve trainers.
4. Begin to build bridges of understanding between the cultures of the museum and the classroom.
5. Promote an attitude of professionalism and accountability
6. Acquire trainers' commitment to use of threaded discussion list software
7. Reinforce the trainers' commitment to use ArtsConnectEd with their students.

In February of 2001 the trainers arrived at the MIA for their first one-and-a-half-day training session. The group of 24 had been split into two groups of 12 to allow for more individual contact time with the Lead Trainer. Each group of 12 would remain together throughout the training in an effort to build community among the trainers. To promote a professional attitude, business cards were made for each trainer and were placed at their seats along with the agenda for the session and the contracts they were required to sign. The following trainers' version of the goals for ArtsConnectEd teacher training was reviewed with the group:

After completing three training sessions in 2001 ArtsConnectEd Trainers will:

1. create an ongoing collaborative laboratory in which criteria for effective classroom use of ArtsConnectEd is discovered, developed, discussed and tested.
2. work as a team to create and evaluate 24 classroom tasks that demonstrate effective classroom use of ArtsConnectEd.
3. increase technological proficiency and aptitude to advocate for increased access to art resources and technology in K-12 classrooms.
4. develop skills that provide leadership in arts education in Minnesota.

In addition, trainers were informed that over the course of their training they would work with the

Fig.3: Trainers install memory in their new computers

Lead Trainer to build a "tool kit" containing skills, materials and resources that they could later draw upon to create effective, high quality workshops for their peers in year two. Once the "house-keeping" was out of the way, the computers were distributed and the rest of the day focused on a series of activities aimed at building overall computer confidence and increasing trainers' familiarity with the specifics of their new tool.

Day two began with breakfast and a question and answer session at the Walker Art Center. Mid-morning, the trainers moved back to the MIA to learn how to configure their computers to access the Internet via both network and dial-up connections. Development of on an on-line ArtsConnectEd tutorial had begun prior to the first training session and the trainers were asked to beta test the program while simultaneously learning about ArtsConnectEd. After lunch the trainers were provided with a field trip behind the scenes at the museum where they followed a work of art from the galleries to the digital photography studio to the computer where the metadata related to the digitized image was input for inclusion in the ArtsConnectEd Web site.

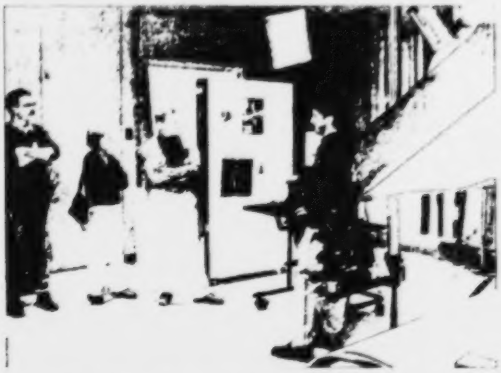

Fig. 4: Trainers visit the MIA's digital photography studio

Beyond providing an in-depth understanding of the internal work process behind ArtsConnectEd, this experience helped cement the trainers' relationship with the museum, making them feel more like museum staff and less like outsiders looking into ArtsConnectEd. Back in the classroom trainers were taught to use ArtsConnectEd's threaded discussion list software (*WebBoard*) and received their assignments for the period until the next training session: use ArtsConnectEd in the classroom with your students, and report to each other about your trials via the threaded discussion list.

Session Two

Goals for Session Two

1. Continue to build bridges of understanding between the cultures of the museum and the classroom.
2. Develop criteria for effective classroom use of ArtsConnectEd based on classroom trials with students.
3. Brainstorm a list of task ideas with trainers based on classroom trials with students.
4. Align task ideas with the Minnesota Graduation Standards using SPACE chart tool.
5. Communicate a procedure for each trainer to write a draft of a task using ArtsConnectEd in a K-12 classroom.

Session two of ArtsConnectEd training began in March of 2001 with a tour of the art storage vaults at the MIA. The MIA Registrar who lead the tour introduced the trainers to the vast holdings of the museum that were not on public view but only available via electronic access. This exercise helped reinforce the unique entrance that ArtsConnectEd provides to the entire collection, well beyond the

Fig. 5: Trainers use Post-It notes to brainstorm task ideas

3-5% of works currently hung on the gallery walls. The Registrar also provided the trainers with information about the lengths that museums go to in order to properly preserve and care for works of art. The experience furthered the goal of immersing the trainers in the mission and culture of museums. Following the tour, each trainer reported on classroom trials with ArtsConnectEd, and several had related student work to share.

The second goal for the trainers: "work as a team to create and evaluate 24 classroom tasks that demonstrate effective classroom use of ArtsConnectEd" was introduced with a definition of a "task." The Lead Trainer had considered that several tasks might be bundled to create a Performance Package, the assessment tool used by the state to determine whether students were meeting Graduation Standards. (Minnesota has since abandoned Performance Packages as its mandated assessment tool, leaving individual school districts to decide on or design tools to measure student achievement.) The task format allowed for the creation of smaller, more manageable units of instruction than traditional lesson plans. A task form was developed and explained so trainers would be familiar with the end product of their goal.

Trainers brainstormed ideas for tasks based on their trials with ArtsConnectEd in the classroom. Once an initial list had been developed, the evaluator led the group to determine criteria for an "effective" classroom task. Trainers discussed and determined the following criteria:

Fig. 6: Trainers review resources in their final training session

1. ArtsConnectEd is required to complete one or more components of the classroom task.
2. Classroom learning outcomes are evident in the task.
3. All tasks are aligned to the Minnesota Graduation Standards.
4. All tasks can be assessed.
5. Use of ArtsConnectEd in the task is planned for a range of available technology.
6. Some tasks are designed to accommodate a diverse range of learners.

Once the criteria were agreed on, the group measured each task idea against the criteria and revised or discarded task ideas that didn't measure up. A final list of all task ideas that met all of the defined criteria was placed on the threaded discussion list after trainers returned to their schools, and each trainer signed up to formally write up one task in the predetermined form. Once all the trainers adopted a task idea to write about, each received a blank task form, a set of written instructions for completing the task form, the SPACE chart that accompanied the task they selected, and a sample completed task form. Trainers were instructed to bring a final task draft to their final training session in June to be critiqued by the group.

Session Three

Goals for Session Three

1. Continue to build bridges of understanding between the cultures of the museum and the classroom.

2. Evaluate 24 classroom tasks that demonstrate effective classroom use of ArtsConnectEd.
3. Demonstrate set-up and use of LCD projector.
4. Define role of ArtsConnectEd trainers in marketing plan.
5. Define procedures for reporting on workshops completed during year two.

Session three of ArtsConnectEd training began in June of 2001 with a tour of the archives at Walker Art Center. The archives house documentation of works of art over the history of the Walker, including artists' models, film, video and audio recordings of performances, correspondence concerning commissioned work or artist residencies, etc. Next, trainers attended a briefing session with the marketing firm hired to promote ArtsConnectEd throughout the state during year two.

That afternoon, the trainers received instructions on care, use and shipping procedures of two LCD projectors purchased for trainers to use in their workshops. A reservation system was set up using the threaded discussion list, where trainers would schedule shipment of the projectors to workshop sites. On day two of the final training session, each trainer distributed copies of the task they had written for group feedback. Following the group discussion, each trainer had two weeks to make final revisions based on the group feedback and submit the final written tasks. Over the summer, an editor edited the tasks and a graphic designer laid out any accompanying handouts. These final versions were compiled in an on-line database that the trainers accessed beginning in the fall of 2001. The trainers could print any of the 24 tasks via the database for distribution in their ArtsConnectEd workshops.

The Evaluation

A three-prong evaluation plan was developed for assessing the effectiveness of training the trainers. First, a pre-survey and a post-survey were developed to measure the growth of the 24 trainers' skills before and after their training. Second, an observer recorded the instructional techniques used and their effectiveness, during each of three training sessions. Third, after each session, the trainers were asked to reflect on and evaluate their training by answering questions on the threaded discussion list. All of the results from each session were compiled, analyzed and shared with the trainers to enable them to make use of what was learned in the evaluation when constructing their own workshops.

Survey Results

The pre- and post-surveys measured quantifiable skills, and the results from those surveys showed that the trainers acquired the skills they needed to successfully transfer their knowledge of ArtsConnectEd to other teachers. Above all, the survey showed that Art Collector, a tool in which users build their own art collections, saw the greatest growth in use over the course of the training period.

Observation Results

The observers at each training session identified the following teaching techniques over the course of the trainers training:

1. Observing a demonstration
2. Clicking along on your own while the instructor provides instructions
3. Group Show and Tell Presentations (trainers share their own experiences)
4. Handouts (paper and electronic)
5. Field trips out of the classroom for a behind-the-scenes museum experience
6. Group discussions
7. Hands-on knowledge and skills gained (as in installing computer memory)
8. Electronic ArtsConnectEd Tutorial
9. Modeling a process
10. Cooperative learning where trainers teach each other
11. Group Critiques
12. Lecture

Results from Questions on the Threaded Discussion List

The list of teaching techniques above was posted on the threaded discussion list and trainers were asked to identify which technique or set of techniques helped them learn the most. They responded that the combination of ALL techniques used was most effective. In response to questions about what worked well during their own training, trainers identified learning computer tricks such as keyboard shortcuts and search tips as one of the ways that their confidence as computer users was increased.

Their confidence as museum users increased when behind-the-scenes information was shared.

Using handouts was also emphasized as a powerful teaching tool, one that addressed different learning styles and provided backup for information given in instruction. In fact when trainers were asked what could have been done to improve their own training experience, the response was "more handouts!" The Education Committee had developed a written and illustrated full-color Teacher's Guide to ArtsConnectEd before the teacher-training program began, and trainers embraced this guide as a basic handout for all their workshops. (Download the guide at http://www.artsconnected.org/classroom)

When the trainers were asked what was learned using ArtsConnectEd in their own classrooms, they reported that students were impressed that the works of art they were seeing on-line could be found in their own state. Once trainers were introduced to ArtsConnectEd in session one of their training, the Lead Trainer began to put scavenger hunts in the threaded discussion lists, requiring trainers to conceive of sophisticated search strategies to solve specific problems.

When asked to pinpoint times when they were most engaged, trainers recalled the periods when they were able to interact with and learn from each other. Making meaning relate to students at a personal level has long been a goal of expert teachers. Tasks and rewards that were meaningful personally as well as professionally were also identified by the trainers as influential teaching tools. Finally, the trainers articulated importance to their overall sense that museum staff and the Lead Trainer treated them with respect and thought of them as professionals. Often the trainers said they strove to behave more professionally because they were treated as professionals.

Task Evaluation

The 24 tasks written by the ArtsConnectEd trainers were evaluated by a team consisting of the evaluator, Lead Trainer, and an education staff member from the Walker and the MIA. The evaluator designed a rubric based on the criteria for effective classroom tasks (see session two), and three members of the evaluation team scored each task.

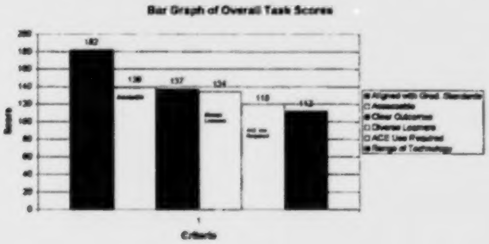

Fig. 7

The resulting scores show that the tasks scored highest in aligning with the Graduation Standards, an exercise that included using a concrete, hands-on tool. The evaluation team was disappointed that the tasks did not score higher on other criteria, and they also expressed the sentiment that overall many of the tasks did not meet museum educators' idea of quality activities. The educators on the evaluation team who had previous experience working with teachers writing museum materials cited similar disappointments in the past.

One theory as to the perceived short-coming of the teacher generated tasks was cultural differences. Classroom teachers are naturally focused on process, switching strategies and changing techniques as a result of the instant feedback they get from students to find "what works." On the other hand, museum educators function in an environment based on academic tradition, where an end product like a research paper or exhibition represents the process that created it. Museum educators respect editorial feedback and go through many revisions and discussions about semantics in order to achieve a final product they can be proud of. Because these two "cultures" are so disparate, the Education Committee wondered if the idea of requiring teachers to achieve museum standards of "quality" was misguided.

Evaluating Trainers' Workshops

The evaluation of the trainers' workshops was conducted via a two-page survey at the end of each workshop. In addition, each trainer was required to reflect in the threaded discussion list on what went well and what could have gone better during each workshop. The evaluator designed a follow-up phone interview for a random sample of those who attended ArtsConnectEd workshops, to be conducted in February of 2002, to determine the effectiveness of the workshops in creating ArtsConnectEd users.

Ongoing Trainer Workshops in Year Two

To date there have been 100 workshops scheduled for the 2001-2002 school year, and 75 of those have already taken place. Many trainers used contacts they had in their areas to schedule workshops during the school district's staff development days, while other trainers took advantage of events already planned such as state conferences that accepted proposals for presentations. Almost all trainers have planned workshops for other staff in the buildings where they teach, and every college in the state that graduates teachers has been contacted with great success with offers of training pre-service teachers. Over 700 people who have attended workshops filled out evaluation forms, and the data has been very positive: 99% say they have acquired the skills and knowledge they need to use ArtsConnectEd, and 73% say they acquired the skills and knowledge they need to use ArtsConnectEd in their classrooms. A full summative analysis of the outcomes of these workshops will be developed in fall of 2002.

Looking Ahead

Museum Fellowships

In the spring of 2002, ArtsConnectEd trainers will have an opportunity to apply for two summer fellowships to develop and complete individual projects related to the goals of ArtsConnectEd and technology use in K-12 art classrooms.

Higher Education Symposium

ArtsConnectEd will also seek to partner with a Minnesota College or University to host a symposium on the use of Web-based resources in K-12 education for all faculties of Minnesota colleges where teachers are educated.

Advanced Regional Workshops

All Minnesota teachers who attended one of the teacher-trainer's ArtsConnectEd workshops during the 2001-2002 school year will be invited to one of three regional ArtsConnectEd workshops in the 2002-2003 school year. Designed by museum staff, these workshops will take ArtsConnectEd users to advanced levels, building on skills learned from the teacher-trainers.

Possibilities for continued or expanded training continue to be explored. While art teachers are perhaps best equipped to understand the value of art museums, there are many potential interdisciplinary applications of ArtsConnectEd. Art teachers require in-depth understanding of other disciplines to be able to teach to teachers of other disciplines. Expansion of teacher training might mean exploring interdisciplinary work with teachers other than art teachers.

Lessons Learned

While the ArtsConnectEd Education Committee is still analyzing the teacher training and formulating what has been learned, some initial lessons can be brought to light. Teachers are in great need of continuing professional development when it comes to using technology in the classroom. ArtsConnectEd trainers were selected for their familiarity with technology; even so, they demonstrated during their training sessions that their knowledge, skills and confidence in working with technology were very diverse. Museums would do well to consider this diversity in skills when developing on-line resources, by designing a range of tools and interfaces for accessing electronic materials.

Giving ArtsConnectEd trainers a look behind the scenes of the museum created dedicated fans who delighted in discovering and exploring all the roles that art museums play. Art museums might have great success making new connections with current visitors and potential visitors by revealing more of the inner workings of the museum to an obviously fascinated public.

Teachers need to be considered professionals, but unfortunately they aren't always treated that way. ArtsConnectEd staff learned that when teachers are treated as professionals, they become more committed, more enthusiastic and more willing to go the extra mile for their students. Supplying teachers with up-to-date materials, professionally produced, and expecting professional behavior created an atmosphere of respect and dedication among the ArtsConnectEd teacher-trainers.

Fig.8: ArtsConnectEd trainers conducting a workshop at the Minnesota Department of Children Families and Learning

Comprehensive online resources, particularly those providing access to the museum collectionz, provide a drastically new paradigm for educators working with art in the classroom. Watching teachers learn to use ArtsConnectEd became like watching teachers and students use a doorway from the school into an art museum. But on-line museum resources like ArtConnectEd can be more flexible for teachers than an art museum. In ArtsConnectEd it is possible to search through storage, organize your own exhibition and write your own labels. Giving teachers museum materials has always been a good idea, but giving teachers tools to access museum collections and all they have to offer provides a more powerful connection.

Finally, as for the train-the-trainer model itself, while the number of teachers and students impacted is impressive in its own right, the long-term impact has yet to be measured. The ultimate goal of sustained and meaningful use of ArtsConnectEd in classrooms across Minnesota as a result of workshops offered by the ArtsConnectEd trainers seems within reach, but documented attainment of that goal is in the future.

Editor's Note: A more extensive version of this paper is included in the electronic edition.

Appendix A - ArtsConnectEd Trainer Contract

This is an agreement between the Walker Art Center acting as fiscal agent for the ArtsConnectEd Project (hereafter known as Walker Art Center) and ______________ (Hereafter known as Contractor) to participate in ArtsConnectEd Training Sessions and following training, to present ArtsConnectEd Minnesota workshops.

This is a two-year agreement beginning February 9, 2001 and continuing through June 30, 2002. The activities of this agreement are coordinated through the ArtsConnectEd Lead Trainer. Failure to fulfill responsibilities outlined in this contract may be considered a breach of contract and may require return of all equipment and stipends and the discontinuation of other contractual benefits.

Contractor Responsibilities:

- Attend three(3) one and one-half day ArtsConnectEd Training Sessions in the Twin Cities:
- February 9/ 10 or 16/17, 2001
- March 23/24 or 30/31, 2001
- June 22/23 or 29/30, 2001
- Subscribe and maintain access to an Internet Service Provider of their choice throughout the two-year term of this agreement February 9, 2001 through June 30, 2002.
- Develop one or more written classroom activities using ArtsConnectEd during the first year of training (February 9, 2001 through June 30, 2001).
- Use ArtsConnectEd with your students in a classroom setting during the first year of training (February 9, 2001 through June 30, 2001).
- Organize and teach a minimum of four (4) Minnesota workshops to at least 15 Minnesota teachers each workshop during the second year (July 1, 2001 through June 30, 2002).
- Participate in evaluation of this program by ArtsConnectEd staff and evaluators hired by ArtsConnectEd

Resources Provided by ArtsConnectEd:

YEAR 1:

1. A Macintosh G3 PowerBook laptop computer.
2. Maintenance and liability for laptops is the responsibility of the Contractor. Walker Art Center will not be responsible for equipment repairs or replacements of hardware or software. All peripheral materials and equipment (such as transportation cases) are the responsibility of the Contractor. Walker Art Center is not responsible for lost or stolen equipment or materials. It is highly recommended that Contractor have coverage from their personal insurance provider.
3. A travel stipend for mileage at $.32 per mile for those traveling over 50 miles one-way to attend the ArtsConnectEd Training Sessions only (February - June 2001).
4. Lodging arranged by ArtsConnectEd staff for one night each for those traveling over 50 miles one-way to attend the ArtsConnectEd Training Sessions only (February - June, 2001).

5. Substitute pay for teachers traveling on Fridays during the school year to attend the ArtsConnectEd Training Sessions only (February - June 2001).
6. A $200 stipend for the completion of each ArtsConnectEd Training Session (February - June 2001).

YEAR 2:

1. A $200 stipend for each Minnesota workshop (up to a maximum of 5 workshops) completed July 2001-June 2002, upon submission of reports including names and addresses of workshop attendees.
2. A maximum of $500 reimbursed for expenses related to running Minnesota workshops completed July 2001-June 2002 (up to a maximum of 5 workshops) upon submission of reports and receipts. Covered expenses include room or equipment rental, copying charges and hospitality costs incurred in four Minnesota workshops. Walker Art Center must approve any other expenses reimbursed.
3. Promotional materials for Minnesota workshops.
4. Limited access to data projectors for use in Minnesota workshops conducted July 2001-June 2002.

Contractor is responsible for all taxes related to stipends and materials, including computer.

Walker Art Center and the Minneapolis Institute of Arts jointly own the copyright for all materials generated by this project. Contractors may reproduce materials for educational purposes only. Any reproduction or distribution for commercial use, fees, or profit is prohibited during or after the terms of this agreement without the written permission of Walker Art Center and the Minneapolis Institute of Arts.

ArtsConnectEd is not liable for any damages or injuries that occur during the course of this program, travel to and from sites related to the program including the Training Sessions and Minnesota workshops.

By signing below I assure that I have read, understand, and agree to the terms of this agreement.

Contractor:

Date:

Social Security Number:

ArtsConnectEd:

Title:

Date:

Appendix B - Trainer Resources and Materials List

Materials

1. Business Cards
2. ArtsConnectEd Stationary
3. Printing Tips handout
4. ArtsConnectEd Tutorial (http://www.artsconnected.org/tutorial)
5. Scavenger Hunts for practice searches
6. Twenty-four classroom tasks
7. Performance Package (assessment)
8. ArtsConnectEd Teacher's Guide
9. Continuing Education Unit Certificate
10. Workshop site technology forms

Support

1. Macinstosh G3 Powerbooks
2. Threaded discussion list (*WebBoard*)
3. $200 stipend per workshop taught
4. LCD Projectors
5. Evaluation forms
6. $500 expense account
7. Museum memberships

Resources

1. SPACE Charts
2. MN Graduation Standards
3. Computer tricks
4. Insider museum stories
5. Museum catalogs
6. Information on Copyright issues
7. Criteria for effective use of ArtsConnectEd
8. Discount Ed. memberships for workshop attendees
9. ArtsConnectEd statistics

Digital Primary Source Materials in the Classroom

Nuala Bennett, and Brenda Trofanenko, University of Illinois - Urbana-Champaign, USA

http://images.library.uiuc.edu/projects/tdc

Abstract

Digital technologies bring museums, libraries, and archives together to enhance learning by providing access to digitized primary and secondary cultural resources along with the more traditional bibliographic materials. At the University of Illinois at Urbana-Champaign, the University Library and the College of Education are developing a collaborative program that integrates digital primary source materials into K-12 curriculum and the educational programs of museums and libraries. *Teaching with Digital Content—Describing, Finding and Using Digital Cultural Heritage Materials* (http://images.library.uiuc.edu/projects/tdc) is a two-year project funded by the Institute of Museum and Library Services. Through this project, we are introducing a broad group of K-12 teachers, museum curators, educators, and librarians to digital cultural heritage materials. Our goal is to provide them with training and professional development activities to enable use of primary source materials in the classroom and in museum and library education programs. In this paper, we will describe the collaborative *Teaching with Digital Content* project and show how the educators are using on-line materials in their learning environments.

Keywords: K-12 teachers, primary sources, education, digital content, classrooms

Background

The *Teaching with Digital Content* project (TDC) project, which started in 2001, built upon previous work undertaken with another project, the Digital Cultural Heritage Community (DCHC) (http://images.library.uiuc.edu/projects/dchc). The DCHC enabled the digitization of materials from East Central Illinois area museums, archives and libraries for 3rd, 4th and 5th grade social science curricula (Bennett & Jones, 2001). This project built upon the concept of a digital community where institutions would contribute content to a database that contained images, text, descriptive information and other multimedia objects to address common themes (Bennett & Sandore, 2001).

The DCHC project aimed to increase the ease with which teachers would utilize these online resources, enabling them to incorporate the new resources in their classroom activities in ways that would be educationally meaningful for their students. We established a framework for the creation of a test database of historical information from area museums, libraries, and archives and tested its viability to meet the curriculum needs of teachers in upper elementary school classrooms in Central Illinois. During the DCHC project, we also sought to develop, document, and disseminate both the processes and products of a model program of cooperation between museums, libraries, and archives.

Several key recommendations resulted from the DCHC project participants (Bennett, Sandore & Pianfetti, 2002). The teachers judged the linking of digitized content to curriculum units and statewide learning standards to be very valuable. The project provided them with an opportunity to match the mandated state learning standards with several classroom activities. The developed database was robust enough to enable very different participant institutions deposit metadata records conforming to the Dublin Core format.

The teachers argued that the quality as opposed to the quantity of resources was important. Educators assigned a high value to the availability of "trustworthy" primary sources via the Web. The time line of the project had not allowed for adequate use and evaluation of digital resources in the classroom. As with many similar projects, the length of time it took to choose artifacts for digitization and the actual digitization took a lot longer than initially anticipated. Finally, the need to increase the ease by which teachers used the images and metadata online in the classroom and for assignments became a significant issue.

BEST COPY AVAILABLE

To further the DCHC project, and to address specifically the resulting benefits and limitations, we proposed and received new funding from the IMLS to bring together ten Illinois and Connecticut museums and libraries and their digital content with K-12 teachers from four school districts in Illinois. This project, administered by the Digital Imaging and Media Technology Initiative of the University of Illinois Library, would also allow us to work more closely with a larger group of museums, libraries and educators.

Museum, Library and Teacher Partnerships

In the new *Teaching with Digital Content* project, K-12 teachers from four Illinois school districts (Bloomington Public School District #87; Champaign School District #4; Springfield Public School District #186; and Urbana School District #116) are working together with ten museum and library partners:

- Chicago Public Library Special Collections Library, Chicago, IL (http://www.chipublib.org/digital/lake/jwedept4_2tn.html);
- Early American Museum, Mahomet, IL (http://node-03.advancenet.net/%7Eearly/);
- Illinois Heritage Association, Champaign, IL (http:// illinoisheritage.prairienet.org/);
- Illinois State Library, Springfield, IL (http://www.cyberdriveillinois.com/library/isl/isl.html);
- Lakeview Museum of Arts and Sciences, Peoria, IL (http://lakeview-museum.org/);
- Lincoln Home National Historic Site, Springfield, IL (http://www.nps.gov/liho);
- McLean County Museum of History, Bloomington, IL (http://www.mchistory.org/);
- Museum of Science and Industry, Chicago, IL (http://www.msichicago.org/);
- Mystic Seaport, Mystic, CT (http://www.mysticseaport.org/);
- UIUC Rare Book and Special Collections Library, Urbana, IL (http://www.library.uiuc.edu/rbx/).

Using an on-line Webboard as the primary means of discussion and communication, the teachers are sharing curricular development and implementation with the museum curators and librarians. Working in conjunction with the College of Education, the teachers have been given guidance on how to develop and engage digital technologies into their knowledge about teaching, teaching about the past and "history," and are learning how to develop and implement these through their lesson plans. These are traditional "lesson plans," in that each includes introductory information about the topic of the lesson, the grade level and the curriculum content area such as social studies, economics or language arts.

Within their lesson plans, the teachers have also identified which of the mandated Illinois K-12 Learning Standards they seek to address with their lesson plans. A recent study of the Education Network of Australia emphasized that "good pedagogy approaches" were necessary to shape the use of digital resources and that children's use of digital technology needed to be guided by its relationship to specific learning goals and integration into specific learning environments (Downes et al, 2000). The Learning Standards are mandated by the Illinois State Board of Education, and Illinois teachers must ensure that lesson plans correspond appropriately. State Goal 16, for example, ensures that students will "understand events, trends, individuals and movements shaping the history of Illinois, the United States and other nations". Within that Goal, there are more detailed categories. The teachers then outline in their lesson plans the activities that meet the Standards and the procedures that the students follow during the lesson. The TDC project differs from the widespread posting of lesson plans on various Web sites, in that those teachers participating in the TDC project suggest specific ways in which the TDC online database might be used in the development of the lesson.

One 4th Grade teacher developed a lesson plan on "Choosing the Lincoln Statue", a competition to design a new statue of Abraham Lincoln for outside the Champaign County Courthouse. During that lesson, students would discuss the chronology of events in the life of Abraham Lincoln, use the TDC database to research the years when Lincoln practiced law in Urbana, Champaign County, search the TDC database for images of Lincoln during the targeted years, and finally, create their own designs for the new statue.

Lesson plans cover a wide range of topics; such as African folk tales, the Civil War, the History of Money, and Settling the Midwest. They also cover a wide span of student ages, 3rd grade through 12th grade. Using the Webboard to share the lesson plans with museum curators and librarians, the teachers also provide the museums and libraries with specific requests for cultural heritage objects for their curriculum planning, development and implementation. Similarly, the museum curators and librarians suggest primary source material from their collections that can be digitized for the teachers' lesson plans. Once artifacts are chosen for digitization, the museum and library participants are developing new metadata which include reference to the Illinois Learning Standards.

Online Database

As discussions develop about the artifacts that would be suitable for the teachers' lesson plans, the ten participating museums and libraries are simultaneously contributing digitized primary source materials and accompanying metadata in Dublin Core format (http://www.dublincore.org) to an online database and search engine. We have made some simple changes to the Dublin Core (DC) field name format used in the database. In particular, we have included fields such as Interpretation, mapped to the DC Description field. (Other fields and their corresponding DC mapping are described in Bennett et al, 2000.) This was due to requests from the teachers for easier identification of field names in the metadata and from discussions among museum and library personnel about how they would like to use the DC for their collections.

As well as the teachers' lesson plans, the Illinois State Board of Education Learning Standards guided and directed the identification and selection of primary source materials for museum curators, archivists and librarians to include in a digital repository . Discussions among the schoolteachers, museum curators, and librarians identified artifacts to be digitized. Although a teacher might have specifically requested artifacts relating to Abraham Lincoln, the museums and libraries, taking into account State Goal 16, might have images of other "individuals and movements shaping the history of Illinois" and suggest those as being usable in the Lincoln lesson plan.

Curators and librarians developed the metadata to correspond to each artifact that would be digitized. The metadata was subsequently formatted in Dublin Core. These artifact images and metadata were then added to the on-line database (http://images.library.uiuc.edu:8081/cgi-bin/htmlclient.exe). The metadata for each artifact image was developed to take into account the different lesson plans and different themes where the image might be used, the different age levels of students and the different requirements of the teachers. For example, one elementary teacher and one high school teacher submitted lesson plans on the two completely different themes of World War II and Women in War (http://images.library.uiuc.edu/projects/tdc/LessonPlans/index.html). The Illinois State Library has a large collection of World War II posters. The participating librarians investigated their poster collection to find some that would fulfill the needs of both teachers and have started digitizing over 100 posters that would help with both lesson plans. In the instance of the "I'll carry mine too!" poster, the library personnel submitted the image and descriptive metadata to the database following several conversations on the Webboard with both teachers. Figure 1a shows the image of a World War II poster held by the Illinois State Library, and Figure 1b is the metadata corresponding to the poster.

Similarly, museum and library partners are digitizing parts of their collections to correspond to other lesson plans. Museums and libraries are assisting teachers by placing the digital objects in their historical contexts. They also identify and help teachers locate and utilize other freely available Web based resources that might be helpful with particular lesson plans. Our experience with the DCHC project was that teachers wanted a trusted source for on-line material to use in the classroom. They do not have the time to research Web resources that would be useful in their lesson plans, but they would be happy to use such resources. Having a trusted group of librarians and museum curators recommend such resources removes that barrier, and the teachers can use other resources in the classroom, or also recommend extra resources to their students. However, the teachers still require much guidance on how to use such resources in the classroom.

Educating the Educators

Besieged by increasing demands for "technological literacy", teachers are invariably being positioned

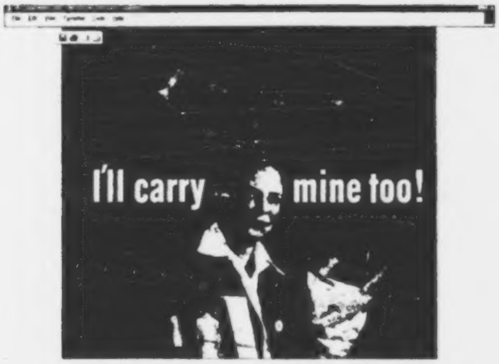

Figure 1a. Poster from the Illinois State Library World War II poster collection – "I'll carry mine too!" *(http://images.library.uiuc.edu:8081/cgi-bin/htmlview.exe)*

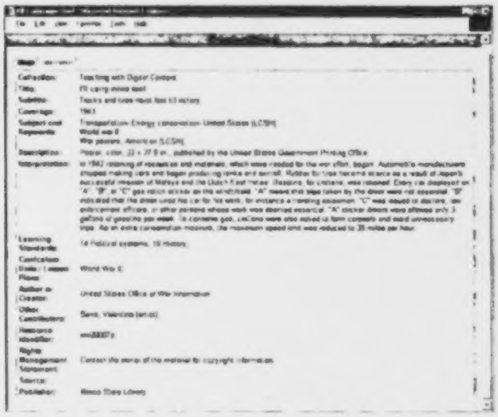

Figure 1b. "I'll carry mine too!" World War II poster descriptive metadata

in an environment where it is essential to appear to be utilizing various technologies. One of the most significant aspects in the contemporary calls for integration of a range of digital technologies into school classrooms points to the "wonders" of computer-mediated learning. New technologies tend to be used in classrooms in ways that are consistent with traditional practices, focusing on goals pertaining to the empirical assessment of student achievement in relation to state learning standards. While teachers in Illinois cannot ignore these standards, the teachers involved in the TDC project are developing lesson plans and activities that are also pedagogically sound and I support student learning rather than being entirely standard directed and driven.

A productive relationship between teachers and museum/library personnel allows teachers to expand their own knowledge base and skills as well as that of their students'

The nature of teaching, the use of digital primary source material, the purposes of the project, and the benefits – and the limits – of technologies in the classroom are all concerns that arise in the various professional development workshops and interactions among and between the teachers, the museum and library individuals, and the project personnel. In presenting practical strategies for effectively engaging with issues about teaching about the past and culture through the use of digital objects, several professional development activities have been planned, instituted, and evaluated, primarily through workshops, summer institutes, and working sessions.

Over the two-year project, the teachers, museum and library participants were introduced to topics specific to teaching with digital content in four workshops. In the first workshop, the teachers were introduced to current research about history education and the tension between knowing what is considered "history" and the past. That the teacher's own historical understanding will influence the students knowing about history, the past, and the tension between the two directed this initial workshop. Teachers were also introduced to the notion of primary sources and the analysis of primary source material (http://images.library.uiuc.edu/projects/tdc/August2001Workshop.htm).

The second workshop brought together individuals from the cultural heritage institutes and the teachers to discuss how the meanings of objects change through digitized imaging, and the interpretive and educational roles the teachers, museum and library personnel have in student understanding. The teachers discussed the role of the object as a medium for learning in the classroom (http://images.library.uiuc.edu/projects/tdc/January2002Workshop_files/). The third and fourth workshops, which will be held at mid-point through the final year of the project, will bring the teachers together again with the museum curators and educators and librarians to further develop learning activities and to engage with issues concerning digital content, notably how differences in the media of

BEST COPY AVAILABLE

primary sources (i.e. photograph, artifact, written text) influence how students learn about the past and how teachers teach about the past.

Assistance is also provided to the teachers to enable them to integrate the digital cultural heritage materials into their curricula. It is not enough for teachers to hear about the theory; they must also be given hands-on assistance to enable them to use the technology in their classroom. Each teacher participant also partakes in *The Moveable Feast* program, a one-week summer institute held at several locations throughout Illinois, which is tailored to the needs of utilizing technology in teaching. This project-based technology institute emphasizes ways to integrate technology in conjunction with lesson plans and the Illinois Learning Standards. As well as the teachers, museum educators and librarians can partake in the Moveable Feast program. Participants receive hands-on training in up-to-date software products and productivity tools that they can use in the classroom. Participants also share ideas for curriculum and Illinois Learning Standards integration.

Since the focus of the project is on understanding, developing, and advancing the social context, or "communities of practice" necessary to engage in educational reform (Lave & Wenger, 1991), the aim is not to participate in producing teachers who have already managed to succeed in traditional ways of teaching history and who now wish to succeed in understanding and addressing the new knowledge and issues associated with the educational intent of digital environments. The aim is to identify and describe the collaborative environments conducive to developing technological and educational competence in the teachers, museum and library personnel, all with an aim of improving the educational experiences of Illinois students. By taking a collaborative approach that focuses on the sociocultural aspects each participating institution brings to education, we hope this project may contribute to a better understanding of collective and interactive attainments, rather than limited individual successes.

Collaboration

The success of the *Teaching with Digital Content* project depends on developing and maintaining a collaborative association among the institutions, the teachers, and the project personnel. Therefore, this project presents numerous pedagogical issues for these individuals. Success depends upon how the pedagogical issues – the differences in the culture of universities, museums and libraries, and schools; the traditional association between museums, libraries, and teachers; and the challenges of educational reform by teachers and museum and library personnel – both bridge the expectations and involvement each individual brings to the project, and promote changes that contribute to student learning.

The purpose of the collaboration is a project that can grow into a learning community (Myers, 1996). While digital technologies figure prominently in this project, social interaction among the participants includes on-line discussions, requests, and in-service events. Reiterative revisions prompt ongoing Web site and curricular development. The term *collaboration* includes two different meanings: to work jointly and to cooperate. For this project, to work jointly requires that the teachers, the museum and library participants, and the project personnel understand each others' expectations, contributions, and knowledge, and work together to ensure that individual and collective intents and expectations are realized.The relationships between teachers and museums and libraries have traditionally been linear, where the teachers accept without question what the museum and library have to offer. In this project, each individual is considered a full and equal partner. On-line and face-to-face discussions bring forward issues and questions about utilizing digital cultural objects, what knowledge is available about the object from the institutions, as well as how the exchange of knowledge between the teachers and students (vis-à-vis lesson plans as a base element) will extend student learning framed within state-based standards.

The tensions between the ideological and pragmatic purposes of digital technologies that influence the project's development, use, and results, which are brought to the project by each individual, prompted the development of a knowledge building community approach (Scardamalia & Bereiter, 1993). To foster communication and collaboration, we needed to move beyond cyberspace to ensure continuous opportunities for face-to-face meetings to address such issues as whether artifacts from museums can be digitized without lesson plans from teachers or whether lesson plans can completely direct the digitization of artifacts and metadata development.

This project speaks directly to educational reform, of re-tooling and changing the ways teachers are engaged in alternative forms of education. Such educational reform requires self-direction and interest on the part of teachers. This self-interest in making an impact on how museums and libraries are used, and how the information obtained from each is then translated into the classroom, is by far the dominant pattern of collaboration among the participants. The image of one's ability to change how students learn, the learning expectations held by the teacher or school community, the museum or library's plans, and the educational system in which each site is embedded, are significant dimensions at both the individual and collective levels. Acknowledging a consistent and on-going commitment to change may result from engaging in the reform process. Those teachers involved in the project need to understand how their involvement changes the teacher's role in educational reform. How successful the implementation and evaluation of the project is involves reconsidering the teacher-curator relationship to provide a working environment where museum and library staff achieve a greater degree of involvement in planning educational experiences and higher levels of expertise in working with digital content. As well, we seek to provide opportunities for the development of technological skills, both for the teachers and the museum and library staff.

To cultivate a collaborative relationship between both groups includes ongoing attempts to implement practices and opportunities for participation by members. The challenge for this project is to work with teachers who must acquire new skills and develop new teaching approaches when working with digital materials. Similarly, museum and library staff must develop innovative ways of presenting access to digitized primary source materials. Both groups are faced with changing expectations and involvement in the project. The teachers become cognizant and critical of digitized artifacts, while museum curators and librarians release the historical hold on interpretation to the teacher for the purpose of student engagement and learning.

Using digitized primary source materials involves fundamental shifts in the philosophies, service methods, and pedagogies of museum curators, librarians, and teachers, specific to particular audiences and intents. Throughout the TDC project, we are seeking to test the effectiveness of introducing several critical components in the process of integrating primary source material into K-12 classroom teaching, including: 1) the use of innovative visual literacy teaching methods in the classroom, 2) a concentrated technology component linking the use of digital technologies with the creation of electronic curriculum materials based on state Learning Standards, 3) easy access to a database containing digitized primary source materials and their descriptive metadata aggregated from a number of diverse museums and libraries, and finally, 4) support (both technical and human) for consistent communication between schools and the cultural heritage associations.

This project binds learning in a context of responsibility and makes each person a contributor to the development of the project. . Thus, the human context surrounding the learning process is at the forefront. For each participant, the learning environment makes possible individual responsibility for learning development, with support from project participants and awareness of being part of a process leading to greater knowledge about utilizing digital content. Cooperation between the participants is key. The scope and complexity of the individual participant's learning curve encourages this cooperation, but it is equally necessary to come up with opportunity and techniques that introduce and reinforce knowledge exchange. It is understood that "to be a partner" means to be in a position to intervene with competence in project development, implementation, and evaluation.

The more each participant is aware of the move towards a collaboration, of the necessity to explore changing relationships between museums and libraries, and changing teacher interaction with each, the more effort each individual may be willing to invest.. From this, the teachers and museum and library personnel can move to uncover the possibilities for their own learning and teaching. This rests on their knowledge about educational reform, which requires new considerations in the management of museum and library artifacts in student learning. Conceptions of history through technology involve the teachers and museum/library staffs' understanding of, knowledge about, and skill development in relation to technologies in their respective field, and their own personal interests and perceived opportunities to develop these interests in relation to TDC project.

Far more critical to the success of the Teaching with Digital Content project appears to be the recognition that social relations frame learning, rather than being derived from it. Each individual will collaborate to refine possibilities for learning. This learning is not limited to either the teachers or the museum and library staff; rather, the learning results when the groups come together to discuss common elements from which to expand.

References

Bennett, N.A., Jones, T. (2001). Building a Web-based Collaborative Database—Does it work? Proceedings, *Museums and the Web 2001*, Seattle, WA http://www.archimuse.com/mw2001/papers/bennett/bennett.html.

Bennett, N.A., Sandore, B. (2001) *Effective Outreach and Collaboration in a Virtual Community:* The Illinois Digital Cultural Heritage Community Project, First Monday (http://firstmonday.org/), July 2001.

Bennett, N.A., Sandore, B., Pianfetti, E. (2002). Illinois Digital Cultural Heritage Community - *Collaborative Interactions Among Libraries, Museums and Elementary Schools*, D-Lib Magazine (http://www.dlib.org/), January 2002, Vol. 8, No. 1.

Bennett, N.A., Sandore, B., Grunden, A.M., Miller, P.L. (2000). *Integration of Primary Resource Materials into Elementary School Curricula*, Proc. Museums and the Web 2000, Minneapolis, MN, pp. 31-38.

Clift, R.T., Thomas, L., Levin, J., & Larson, A. (1996). *Learning in Two Contexts: Field and University Influences on the Role of Telecommunications in Teacher Education*, Presentation to the American Educational Research Association.

Downes, T., Arthur, L., Beecher, B., Kemp, L. (2000). *Appropriate EdNA Services for Children Eight Years and Younger*. Report commissioned by the Education Network Australia On-line Pathways Project, and Education Limited.

Fullan, M. & Stieegelbauer, S. (1991). *The new meaning of educational change*. New York: Teachers College Press.

Myers, C.B. (1996). *University-school collaborations: A need to reconceptualize schools as professional learning communities instead of partnerships*. Paper presented at the annual meeting of the American Educational Research Association, New York.

Hargreaves, A. (1995). *Changing teachers, changing times. Teachers' work and culture in the postmodern age*. New York: Teachers College Press (Professional Development and Practice Series).

Lave, J., Wenger, E. (1991). *Situated learning: Legitimate peripheral participation*. New York: Cambridge University Press.

Scardamalia, M., Bereiter, C. (1996). *Engaging Students in a Knowledge Society. Educational Leadership*. Alexandria, Virginia: Association for Supervision and Curriculum Development. Vol. 54, No. 3, November, 6-10.

Oakes, J., Hunter Quartz, K. (1995). Creating New Educational Communities. *Ninety-fourth yearbook of the national society for the study of education*, Part I. Chicago, Illinois: University of Chicago Press.

Wells, G. et al. (1994). *Changing schools from within: Creating communities of inquiries*. Toronto, Ontario: OISE Press; Portsmouth, New Hampshire : Heinemann.

Statistics, Structures & Satisfied Customers: Using Web Log Data to Improve Site Performance

Darren Peacock, National Museum of Australia, Australia

Abstract

Qualitative and quantitative evaluation of visitor experience in museums has a proud and well established tradition. Long before "customer relationship management" became a ubiquitous catch cry, museums were engaged in rigorous and sophisticated analyses of their audiences. Similarly, museums have also been at the forefront of developments in online content delivery. Yet, the culture of rigorous evaluation applied to traditional visitor research is not nearly so apparent in the online museum environment. As competition amongst online content providers becomes more intense, museums need to embrace a more rigorous approach to understanding and developing their web audiences. Building on our established traditions of audience evaluation, museums can once again lead the way in developing understandings of how visitors explore and engage with content in the new realm of virtual experience.

Web log analysis is an under-utilised approach to understanding and testing visitor behaviour on the web. Every visit to a site leaves a potentially rich vein of information for any willing data miner. Utilising that data effectively to understand the visitor's experience is essential to building web sites that work. The National Museum of Australia is using the analysis of web log data to inform the redesign of its online presence. Using new analysis tools, historical log data is being mined to test hypotheses about user behaviour and to develop new approaches to site structure and design. As the new site is implemented web log data will be used as the basis for the ongoing study of changing patterns of visitor behavior.

Key words: evaluation, log, museum, statistics, web, pathing, metrics

Introduction

Why should we evaluate websites? Like any museum endeavour, if we start out without a map of what we hope to achieve, it is likely we'll end up some place quite different. Evaluation helps us find our bearings and to move on to the next stage.

Another answer to the question 'why evaluate' is that, unless we do, future developments will rely on guesswork rather than research. In the online world, it is all too easy to let the technology drive the agenda. Evaluation pulls us back to the world of the user and commits us to an active engagement with our online visitors and their needs, expectations and experiences.

By now, most museums are through the establishment phase of their online evolution. The challenge now is to establish a sustainable program of continuous improvement in online content, based on what we know about the medium and our visitors' use of it. Museums are well versed in the tools and techniques of evaluation. In the past four decades, visitor studies have fundamentally altered approaches to the development and delivery of core museum services. As Hein (1998) observes, "interest in visitor studies expanded dramatically in the 1960's, coincident with both increased government spending on a wide range of social services and increased application of social science research to examine these activities."

Exhibitions, activity and education programs and the publics who use them have benefited enormously from the application of rigorous techniques of critical evaluation. As museum spending for online service delivery increases, the need for audience research also grows, both to justify further investment and to substantiate returns on current investments.

Online content delivery is here to stay as a core function of museums. What then will we use as our guide in developing online material and creating and extending our audiences? What kind of evaluation tools can we use? Which will yield the best results in terms of improving the user experience? How can we establish a cycle of continuous improvement as online audiences grow and their expectations increase? These are the questions this paper addresses.

BEST COPY AVAILABLE

Museums and the Web 2002: Proceedings
Archives & Museum Informatics, 2002, p. 157

Museums were quick to seize the opportunities offered by the web. Now, just having a site is no longer sufficient. Having a site that demonstrably meets and develops user expectations is essential. As the number of surfing options grows exponentially, museum sites have to compete for the eyes and mouse clicks of users in an increasingly cluttered and competitive cyberspace.

Any evaluation approach requires both data and a methodology, or framework of analysis. This paper explores some of the ways in which the National Museum of Australia is using web analysis tools to shape its future directions in the delivery of online services.

In particular, it explores the potential of quantitative analysis, based on web server log data, to convert these ephemeral traces of user experience into a strategic management approach for online service delivery. My goal is to present a methodology and a set of potential e-metrics for evaluating and improving user experience on museum websites. In this model, customer satisfaction, measured through quantitative analysis, provides benchmarks for site performance and directions for future development.

National Museum of Australia

The National Museum of Australia is in the unusual position of having established an online presence before it became a physical reality. The museum has operated online since 1995 (www.nma.gov.au). However, it was only a year ago that the museum opened its major visiting facility on the shores of Lake Burley Griffin in Canberra, the nation's capital.

The museum building was a flagship project celebrating the Centenary of Federation, when, in 1901, six British colonies united to form the nation of Australia. The National Museum of Australia benefits from having being born digital. It is a truly wired museum and incorporates a full television and radio broadcasting facility, which has already been harnessed for producing webcast content.

From the very beginning, it was understood that virtual audiences- both broadcast and on line- would be integral to the National Museum's operations. Canberra is a city of just 300 000 people, relatively distant from the main population centres of Sydney and Melbourne. To serve the whole of the nation, the National Museum has a major commitment to using technology to reach diverse and remote audiences. Moreover, as an institution, we are positioned as a forum, not a temple- to use the terms of Duncan Cameron's 1971 thesis- committed to stimulating, convening and contributing to national debate. Broadcasts, webcasts and other electronic outreach are vital to our mission and reason for being. Building our online presence is a key strategic priority.

New medium, new audiences?

When we look for models of website evaluation, we can turn to the museum world's own rich repertoire of visitor research tools and techniques and also to the market research paradigms of commercial online services. Our choice of method raises the question of how online museum audiences differ from physical museum visitors. Do they have more in common with gamers and e-shoppers? Are their expectations of online museums shaped by their experience of traditional museums or by their experiences in the online world? What models and benchmarks of customer service can we or should we embrace from non-museum sites?

This takes us into the territory- well trodden by visitor studies research- of who do or should museums serve? In the online environment, this becomes a question of who are our audiences (users), how do we attract them and how do we serve them? What models of customer service do we wish to establish for online museum services? How will we know if we are meeting users' needs? And just what are their needs?

Despite the established tradition of evaluation research within museums, models of museum website evaluation have been slow to emerge. Studies of individual sites or comparative analyses have shed some light on user expectations and experiences. A number of papers previously presented at this forum-Bowen (1997), Chadwick and Boverie (1999) and Semper and Jackson (2000)- provide much food for thought.

The results from these studies may answer specific questions about particular site content and design, but are yet to establish a generalised hypothesis about the nature of user experience on museum sites. Teather (1998) observed that 'user study is seldom taken up as it is seen as too expensive….'

Perhaps rather than cost, the problem lies in deciding what to test about user experience. Certainly there have been extensive efforts in the area of useability testing, where functions, navigation and graphic design can be tested by potential users. But usability testing is not the same as usefulness testing. Without an understanding of user needs and expectations, a site may be well designed without being useful.

Not surprisingly, much research attention to date has been focussed on the sites themselves, rather than directly on audiences.As web technologies have rapidly evolved, there has been much experimentation with the form and content of museum sites.

Many sites have evolved from the first generation brochureware sort, to become multi-dimensional interactive spaces, including databases and complex multimedia presentations. An initial focus on the potential of technology has been a necessary starting point, but organisational goals are now developing beyond just having a site to having an online presence which is driven by the strategic mission and goals of the organisation, not just the latest technological innovations.

To ensure that our online offerings are relevant and useful, we must proceed from an audience (user) perspective. Our analysis of online services should encompass the whole visit experience, from its genesis to its conclusion, taking customer satisfaction as its goal. This approach has already been advocated effectively by Falk and Dierking (1992, 2000) in their research methodologies for museum visitor studies.

Obviously, any visitor research has to compete for funding and attention against other more pressing operational needs, including traditional forms of visitor research. Yet web evaluation is in fact potentially cheaper than other traditional forms of observational and survey research. Moreover, research into web audiences can extend and enrich our understandings of visitors who do come through the door.

Research choices

Research into visitor experience comes in many forms, typically dividing into qualitative and quantitative methods. Hein (1998) has offered a useful summary of popular techniques.

Front-end usability research techniques have used a range of methods for assessing the interaction between people and computer interfaces. Usability research for human-computer interaction has a long history. People like Jakob Neilsen have made a major contribution to our understanding of these interactions.

Yet, as Teather (1998) observed, usability testing approaches need to be more closely aligned to models of museum visitor research. Usability testing may tend to centre too much on engineering and technical design, focusing on the site or interface itself, rather than putting the visitor at the centre of the analysis.

To get to the heart of usability, we need to understand users' needs and motivations, not just their responses. We need to model those needs and to design and test accordingly, not just for functional effectiveness and efficiency, but for customer satisfaction.

This is where commercial market research approaches can help us with the task of differentiating our users and their needs. Techniques of audience segmentation and the personalisation of content begin to recognise the diversity of users and user needs, rather than focussing on the site itself. Recent publications by Inan (2002), Sterne (2001), and Grossnickle and Raskin (2001) provide excellent overviews and critiques of research and evaluation methods being developed in the commercial online world.

Let the servers do the surveying

But what data shall we use for our analysis? Log file data, quietly accumulating on web servers, is the cheapest and least exploited data source for understanding museum web visitors and their needs. No other survey technique generates as much data for so little effort. The trick is to turn it into useful information and practical applications.

Unfortunately, due to an emphasis on its deficiencies, log data has a rather poor reputation. I would like to demonstrate how, for me, it sits at the foundation of our approach.

Log data has been dismissed as useless and inaccurate because of some inbuilt limitations which arise from the way the web works. After initial early enthusiasm for web log analysis, in recent years, commentators have tended to pay lip service to the data and move on to suggest other alternatives. Analysis of museum web sites has typically focussed on other qualitative forms of analysis gathered through surveys of users or by observational analysis.

I will recap briefly the nature and limitations of log data for those who are unfamiliar with or have been deterred from further exploration. I would then like to examine how, notwithstanding those limitations, log data may be used as the basis for a methodology to test visitor satisfaction with their online experiences.

Getting beyond the traffic

Web logs were initially used as traffic counters. Fundamentally, they are a tool of the technicians. They were originally set up to measure the volume of page requests, which provided important information to help plan server capacity. The widespread, erroneous use of 'hits' as a measure of site effectiveness contributed significantly to the discrediting of log based research.

The exploration of log data as a potential source of market intelligence by some coincided with the discovery of its deficiencies as a measurement tool by others. Because of the way the web operates, in particular the process of caching, logs do not reliably count the total number of page requests or user sessions. However, this does not, as some have suggested, render log file data as "worse than useless". With caution and appropriate caveats, log data is still a rich and useful source of relevant information about the user experience.

If we move beyond simply tracking site traffic, we can reconceptualise log data as a survey sample of web visitors. It is not the whole story, but a significant and valuable sample of the whole.

Like all sampling techniques, log data has inherent biases and blind spots. As any visitor researcher will attest, the perfect sample of any visitor population is hard to find. Yet if we acknowledge the deficiencies of the log derived sample of web visitors, we can possibly turn them to advantage.

The biases of log data are at least constant and predictable. Because of caching it consistently underreports repeat visitors and users from the most popular ISPs. Log files may also inflate the number of unique visitors to a site, as the same user may be logged with multiple IP addresses during a single session.

These skews in the data suggest that the sample users recorded in the logs are more likely to be first time visitors to the site. In a study of customer satisfaction it is these users who are probably of more interest than repeat visitors. Arguably, we may consider that repeat visitors were at least partly satisfied during their previous visits to motivate a return.

Log derived data about user behaviour may therefore be most revealing in ascertaining how new visitors access and make use of the site and whether they leave satisfied by the experience. The data recorded in the logs about these users can provide useful insights into the initial impressions and motivations of first time visitors to our site. After all, in an increasingly competitive cyberspace, new visitors are the main hope for growing online museum audiences.

The advantages of log files are that they produce quantitative data that can be subjected to statistical analysis. The data samples are large and can be tracked over time. The data produced is a record of actual user behaviour rather than reported or assumed activity. Log data is recorded free of observer or questioner bias.

If we accept that, despite its inadequacies, log data is the most comprehensive source of data about online visitor behaviour, how do we build a model to apply that data to the measurement of site effectiveness, that is, of customer satisfaction?

Modeling the user experience

In modeling the user experience, we should return to the concept of visitor needs. Like visitors to museums generally, web visitors are not empty vessels waiting to be filled with museum content. Falk and Dierking (1992) have elsewhere made the case for acknowledging 'visitor agendas' amongst physical museum visitors. In the online realm as well, visitors often have clear, conscious agendas of their own and sophisticated searching skills to pursue their goals. If those agendas are not satisfied, visitors are soon lost and may never return.

Psychologist Abraham Maslow offered one of the simplest and most compelling frameworks for understanding human motivation with his Hierarchy of Human Needs (1954). Adapting the idea of Maslow's hierarchy to the needs of online museum visitors, I would like to suggest a framework for testing user experience based on analysis of web log data.

The four tiers of the proposed framework (Figure 1) map the stages by which users access and explore a site. At each level there are a set of log diagnostics which can be used to measure the pathways and obstacles to user satisfaction. Together, these diagnostics form a set of e-metrics which can be applied across institutions and across time to benchmark site effectiveness.

Level 1: Can I find it?

At the lowest level of the hierarchy, we are concerned with the most fundamental issue- how the visitor gets to the site. Log data enables us to monitor and examine the ways in which site traffic is generated.

This is a fundamental diagnostic, showing the effectiveness of search engine registrations, links from other sites and our own site promotions. For example, on our site, eight of the top ten external referrers in the month of January 2002 were popular search engines. The other two were links established by the museum with other organisations as cross-promotions, one a local tourism centre, the other a scientific research institution.

This level also shows the context of the visit. Most users will visit several sites in a single internet session. There may be a logical sequence from one search or site to another. Analysis of the search criteria used to locate the site will show what search terms visitors are using to find the site.

From the first level of the analysis framework, we can establish a picture of our place on the information superhighway and the routes traversed to locate our sites. Comparison between sites enables us to establish a sense of where and how the site's profile can be enhanced.

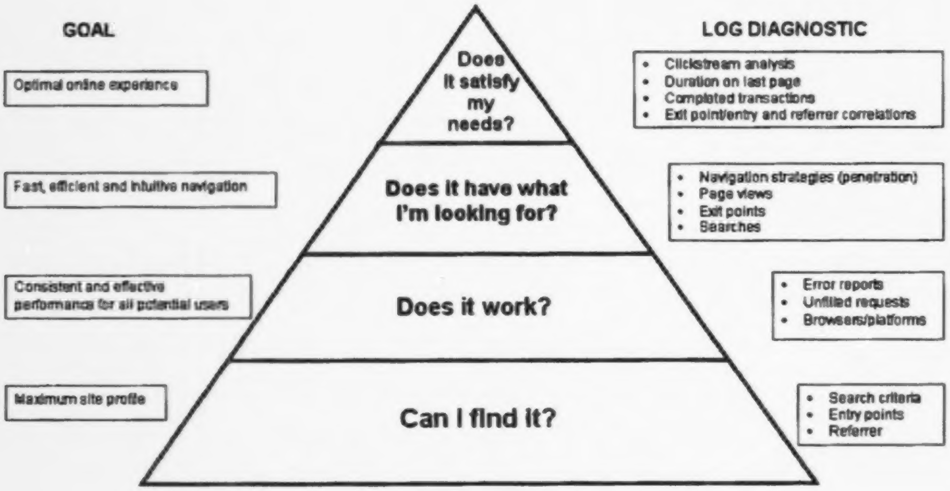
GOAL
Optimal online experience
Fast, efficient and intuitive navigation
Consistent and effective performance for all potential users
Maximum site profile
Does it satisfy my needs?
Does it have what I'm looking for?
Does it work?
Can I find it?
LOG DIAGNOSTIC
Clickstream analysis
Duration on last page
Completed transactions
Exit point/entry and referrer correlations
Navigation strategies (penetration)
Page views
Exit points
Searches
Error reports
Unfilled requests
Browsers/platforms
Search criteria
Entry points
Referrer

Figure 1. Hierarchy of web user needs

BEST COPY AVAILABLE

Some of the key questions addressed at this level include:

- How many people come directly to the site home page?
- What proportion arrive through search engines?
- What proportion arrive at other parts of the site?

At the National Museum of Australia, we discovered from log data that many visitors were coming to the site to plan a trip to the new museum. Information about daily events at the museum was an obvious goal for a high number of visitors to the site. To make it easier to achieve the goal we created a "What's On Today" link in a prominent position on the home page. The log data shows a clear response to this innovation. In the week we introduced the link, some 200 visitors clicked there first. After a month, it had become the second most popular link from the front page.

Level Two: Does it work?

The second level of the needs hierarchy examines the user experience from the practical perspective of site performance. A museum perspective on site performance might encompass issues about availability and the management of traffic volumes. From the user point of view, the key issues are speed and reliability.

The goal of performance standards for this level in the hierarchy relate to the fast and effective delivery of pages to the broadest range of potential users, regardless of their operating systems and network connections. Log data provides effective monitoring of the content delivery experience of users. Error logs show such defects as broken links, server errors and refused requests.

Unfilled requests showing aborted page requests provide evidence of user frustration with slow downloads. Recurring patterns of unfilled requests point to problems with design and file sizing. Fortunately, I can report that >98% of the page requests for our site in January 2002 were delivered successfully (Status code 200).

Users come in all shapes and sizes. They arrive at our sites using a bewildering array of hardware and software. Profiles derived from log files of the browser and operating software deployed by our users sheds light on visitors' needs and helps establish minimum and maximum standards.

For example if you discover- as we did- that 60% of your visitors are using Internet Explorer v5 and above to access the site, this may be the median (or the baseline) for which you design. You can also accurately identify the extent to which you risk disenfranchising potential users by pitching your site design beyond the capability of their software or network connections. These measures can also be compared to industry standards to assess the visitor population of the site against the whole population of internet users.

For the month of January 2002, the distribution of browsers employed by visitors to the National Museum site were as follows:

Browser	% of sessions
Internet Explorer	63.3
Netscape	25.3
Others	11.4
TOTAL	100

Version of IE	% of IE sessions
5.5	35.5
5.0	24.1
6.0	17.4
5.01	16.0
4.01	4.0
Others	3.0
TOTAL	100

Figure 2. Breakdown of visitors by browser type and version, National Museum of Australia website, January 2002

Level 3: Does it have what I'm looking for?

The third level in the hierarchy moves from the operational accessibility of the site to the effectiveness of the content and its organisation.

This level is concerned with navigation and how it facilitates or impedes the user experience. Log data in respect of this area requires more complex interpretive approaches. The data can be used to test hypotheses about the search goals and strategies of visitors to highlight potential areas for improvement.

At a more fundamental level, pathing analysis helps us to understand the motivations of visitors. If we analyse where they enter a site, how they move through it and where they exit, we can test hypotheses about their motivations and goals.

This type of material proves most potent in understanding how visitors interact with your site. It also serves to remind us that many visitors merely pass through our sites on their way from and to somewhere else. The trick is to hold their attention. If we fail to capture their interest beyond the click that brought them to our sites, there is something seriously wrong. The analogy is the physical visitor who enters the building and departs without looking further. Certainly, this counts as a visit, but we can hardly consider it a satisfied customer.

Our first exploration of pathway analysis was revealing and sobering. Our web analysis tool- *Webhound* – from the SAS Institute, generates a top ten pathway report for any given period. This report has shown a potential problem with a very high level of attrition from the home page, which we are hoping to address with a redesign.

The January report- before the redesign- shows the following as the most popular 'triplet' combinations (sequence of three pages) on the museum site:

Rank	Sequence
1	Entry/Home/Exit
2	Entry/Home/Visiting the museum
3	Entry/Home/Exhibitions
4	Entry/Home/What's On
5	Entry/Home/Jobs
6	Home/Jobs/Exit
7	Home/Visiting/Opening Hours
8	Entry/Home/Online showcase
9	Entry/Featured Exhibition/Exit
10	Home/Visiting/What's On

Figure 3. Most popular three page sequences in rank order, National Museum of Australia website, January 2002.

Webhound also maps the sequences dynamically so you can assess the popularity or otherwise of any particular pathway through the site. This is also an excellent way to track the level of penetration into the site and the extent of decay with each click point.

Level four: Does it satisfy my needs?

Atop the pyramid is the holy grail of customer satisfaction. How do we know that we've met the needs of visitors? Few, if any, give us direct feedback; repeat visitation is another form of vindication- if it can be tracked; ad hoc surveys may establish some measure of satisfaction.

Once again it is log data that provides the largest, most objective and most valuable record of user experience. That is, if we know how to interpret it. This model of user needs proceeds from the assumption that visitors to museum sites are goal directed. Whether they are browsers or searchers, once they have clicked into the site, they are demonstrating a need or an interest which they believe, rightly or wrongly, that the site can fulfil.

How they pursue that need or interest and the extent to which it is fulfilled, can only be accurately tracked by reference to the logs. A number of diagnostics can be devised to measure satisfaction in the sense of a goal that has been achieved. I suggest four as the basis for testing hypotheses. Two of these relate to the last thing that a visitor does before leaving a site. For example, the last page visited is revealing. For the month of January, 2002, on the National Museum of Australia site, they were as follows:

Rank	Page
1	Home
2	Jobs
3	What's On
4	Feature exhibition
5	Exhibitions
6	Tenders
7	Visiting the museum

Figure 4. Last page visited in rank order, National Museum of Australia website, January 2002

Once again, it suggests that our visitors have clear goals for their visits. Correlation of the exit page with the original referrer or search terms can be used to confirm the intentionality of the visitor and

162

INDICATOR	WEB LOG MEASURE	MEASURE AS
TRAFFIC	Sessions	Number
	Page views	Average number per session
	Duration	Average time per session
	Domain	% by location/type
	Time of visit	Spread by day/time
SATISFACTION		
(Level 1)	Entry point	Top 10 by volume
	Entry search	Index of search terms used
	Referrer	Top 10 by volume
(Level 2)	Errors	Number by status code
	Unfilled requests	% of total
	Browser	% by type/version
	Platform	% by type/version
(Level 3)	Pathing	Top 10 sequences
	Penetration	% reaching each level
	Exit point	Top 10 by volume
	Internal searches	Index of search terms used
(Level 4)	Exit point by referrer	Correlation
	Duration on last page	Average time
	Exit point by entry point	Correlation
	Completed transactions	% of forms submitted /accessed

Figure 5. 20 log based metrics for evaluating museum websites

to assess if their goals appear to have been achieved, eg. to find a job, a tender, or 'what's on' information.

A set of log-based diagnostics

From this framework therefore, I would like to propose 20 log based diagnostics for evaluating the performance of museum sites (Figure 5). The first five, which are traditional 'traffic' indicators, showing the volume and origin of page requests: sessions, page views, duration, domain, time of visit.

The second group of 'satisfaction' indicators relate to the four levels of the user needs hierarchy outlined in Figure 1. As the National Museum of Australia extends its online presence, these are the metrics we will use to track performance and to inform design and content development decisions. Imperfect science perhaps, but an improvement on the guesswork of the past.

Conclusion

Museums have long been adept at analysing their customers through rigorous and diverse analysis of visitor needs and behaviour. The internet offers museums vast new ways of reaching and engaging with audiences, both existing and those who have never visited before. We have to learn how to apply new tools of analysis with the same rigour and spirit of open enquiry that has characterised the last forty years of museum visitor research.

With our already established practices of research, we should be able to emulate the best practice evaluation models emerging in the commercial online world.

This paper has suggested a possible starting point for building an evaluation model for online museum services. Web log data, despite its limitations, is an essential part of that framework. It can be supplemented by other research, but it offers unique advantages as a source of data on the user experience.

The implementation of such frameworks will ensure that our online offerings are dynamic in response to customer needs, not just technological innovation. In an increasingly competitive cyberspace, building and nurturing customer relationships will be essential. We've only just begun this new voyage of audience discovery. I hope the map offered here helps.

References

Bowen, J. (1997) The Virtual Library museums page (VLmp): Whence and Whither? In D. Bearman and J. Trant (Eds.) *Museums and the Web 97: Selected Papers.* Pittsburgh: Archives and Museum Informatics. 9-25

Cameron, D. (1971) "Museums, Temple or Forum", *Curator* 14 (1).

Chadwick, J. & Boverie, P. (1999) A Survey of Characteristics and Patterns of Behaviour in Visitors to a Museum Web Site. In D. Bearman and J. Trant (Eds.) *Museums and the Web 1999, Selected papers from an International Conference.* Pittsburgh: Archives and Museum Informatics. 163-172

Falk, J.H., Dierking, L.D. (1992) *The Museum Experience.* Washington DC: Whalesback Books.

Falk, J.H., Dierking, L.D. (2000) *Learning from Museums.* Walnut Creek, CA, Alta Mira Press.

Grossnickle, J. & Raskin, O. (2001) The Handbook of Online Marketing Research: Knowing Your Customer Using the Net. New York: McGraw Hill.

Harms, I. &, Schweibenz, W. (2001) Evaluating the Usability of a Museum Web Site. In D. Bearman and J. Trant (Eds.) *Museums and the Web 2001, Selected papers from an International Conference.* Pittsburgh: Archives and Museum Informatics Museums and the Web 2001 43-52

Hein, G.E., (1998) *Learning in the Museum.* London: Routledge.

Inan, H. (2002) *Measuring the success of your website: A customer-centric approach to website management.* Sydney: Pearson Education Australia.

Maslow, A. (1954) *Motivation and Personality.* 3rd edition 1987, New York: Harper and Row.

Semper, R. & Jackson, R. (2000) Who's Out There. In D. Bearman and J. Trant (Eds.) *Museums and the Web 2000, Selected papers from an International Conference.* Pittsburgh: Archives and Museum Informatics. 179-197

Sterne, J. (2001). *World Wide Web Marketing: Integrating the Web into your Marketing Strategy.,* 3rd edition, New York: Wiley.

Teather, L. (1998) A museum is a museum is a museum or is it: Exploring museology and the web. *Museums and the Web 98 papers* Pittsburgh: Archives and Museum Informatics. http://www.archimuse.com/mw98/papers/teather/teather_paper.html

How Do You Like To Learn? Comparing User Preferences and Visit Length of Educational Web Sites

David T. Schaller and Steven Allison-Bunnell, Educational Web Adventures, Minda Borun and Margaret B. Chambers, Museum Solutions, USA

Abstract

Developing effective public education sites for the Web requires an understanding of both learning theory and what appeals to learners. A recent study commissioned by IBM found that Web learners prefer passive entertainment experiences to more demanding interactive experiences (Karat et al, 2001). If people learn best in active modes, but prefer passive Web experiences, how can we develop sound educational activities that attract and appeal to a broad audience?

This paper reports results of a study designed to determine people's preferences for different types of Web-based educational activity. The primary research question was: How do people's preferences vary among types of Web-based learning activity? We identified six activity types for comparison: Creative Play, Guided Tour, Interactive Reference, Puzzle/Interactive Mystery, Role-playing Story, and Simulation. A team of Web developers who work with museums and other learning sites collaborated with a team of educational researchers who work primarily with museums to conduct a survey of visitors to five different types of educational Web site. Two kinds of data were collected: 1. User exit surveys, eliciting an evaluation of the study site and preferred genre or type of learning activity, and 2. Server statistics indicating the duration of stay. Results indicate that there are clear differences in the type of Web-based learning activity preferred by adults and children. Adults are more likely to select Interactive Reference or Simulation whereas children prefer Creative Play and Role-playing Stories. The adult sites yield more straightforward cognitive information while the sites preferred by children allow more personal choice and interaction. Apparently, adults bring an intrinsic motivation to the learning experience. They know what they want to learn and they want to learn it in the most direct way. Children, on the other hand, need to be motivated. They respond positively to the opportunity for interaction and choice within a goal-based environment that offers them an extrinsic purpose.

Keywords: learning preferences, learning theory, Web-based education, goal-based scenario, intrinsic motivation.

Applying Learning Theory to Interactive Media

In the past decade, the Web has grown from a text-only tool of academia to a dazzling universe of ideas, community, commerce, and vanity, with a corresponding increase in its multimedia capabilities. How can the Web best be used for education? Applying learning theory to an immature medium like the Web is challenging, but several basic criteria for learning can safely be applied:

- "Education is not an affair of 'telling' and being told, but an active construction process"(Dewey, 1923).
- "Learners do not learn directly from technology [or teachers, or books]; they learn from thinking about what they are doing" (Jonassen, 1999).
- Learners must be motivated, which requires an "emotional connection, challenge, and payoff" (Healy, 1994).

Somewhat more controversial is a key tenet of constructivism: "A range of results are possible *and acceptable*" (Hein, 1998). This tolerance for divergent outcomes distinguishes constructivism from discovery learning, in which "by engaging learners in activity...they will arrive at the *correct conclusions*" (Hein, 1998). Constructivism suggests that learning activities should allow multiple outcomes, each of which need only "'make sense' within the constructed reality of the learner" (Hein, 1998). (For additional analysis of theories of learning applied to the Web, see Schaller and Allison-Bunnell, 2001)

BEST COPY AVAILABLE

Beyond pedagogical approach, museums and other organizations devoted to leisure learning must decide on the desired type of educational experience. Gammon (2001) offers a useful typology:

- **Cognitive:** Acquire and assimilate new knowledge into existing schemas, apply existing knowledge, connect concepts, draw analogies.
- **Affective:** Challenge beliefs and values, appreciate viewpoints in other people, inspire interest, curiosity, awe and wonder, associate curiosity and thinking with enjoyable experiences.
- **Social:** Develop skills of co-operation and communication.
- **Developing skills** (mental and physical): Prediction, deduction, problem-solving, investigation, observation, measuring, classification, testing theories, making and telling stories, decision-making, manual dexterity, craft skills, etc.
- **Personal**: Increasing self-confidence and self-efficacy; motivating to investigate further.

With these issues in mind, how has the Web fared? The Web is a form of interactive multimedia, or IMM. Some educational researchers and practitioners praise IMM's ability to use audio, video, text, and immersive environments to appeal to multiple intelligences (Veenema and Gardner, 1996). Others see in IMM the chance to move beyond passive learning modes and engage students in more active learning experiences (Prensky, 2001; Crawford, 1982; Viadero, 1996; Tipping and Graesser, 1996; Bearman, 1997; Plowman, 1996b).

However, actual evaluation of IMM products has shown that logistical problems often get in the way of fully realizing IMM's potential. Many novice learners have found the navigational choices offered in IMM programs bewildering. In studies of classroom use of IMM, students have needed considerable teacher assistance to make use of the programs. (Veenema and Gardner, 1996; Plowman, 1996b; Bearman, 1997). The problem lies in the very freedom afforded by IMM's non-linear structure: "Being a user-controlled medium, the learner expects to have control, and yet a learner does not know enough to be given full control" (Laurillard, 1996). Novice learners need more guidance and structure to ensure that they find content that is both engaging and appropriate to their knowledge level.

While classroom teachers with sufficient time, skill, and motivation can overcome these difficulties and provide the necessary guidance to make use of the experience offered by IMM, this is not an option for Web-based leisure learning experiences. Web sites must attract an audience and create a self-contained experience that is satisfying and hopefully educational. Thus we must account for what people want as well as how they might learn.

A recent study conducted by IBM suggests that, given a choice, leisure learners seek relief from bewildering interactive software. This formative research revealed that:

> most participants did not express interest in Web sites that involved active interaction with the content or other people. [They strongly preferred being] guided through an experience or discovery process (Karat et. al, 2001).

Some participants in the IBM study "viewed the more interactive design concepts and existing Web sites as work, not entertainment" (Karat et. al, 2001). Indeed, the learning modes that IBM researchers offered participants were either quite passive (Guided Tour) or quite active (a searchable database of images and information, a chat room, and an online journal). The latter may engage the devotee or a student doing a research report, but can easily overwhelm those who lack an existing interest in the subject and the intrinsic motivation to explore it.

Based on the results of their study, IBM developed a site featuring on-line tours, hosted by curators and other experts, and delivered via streaming video—essentially a TV-like experience with links to additional information. Summative evaluation of the site found that "users interacted relatively infrequently with the [on-line] tours, and the less they interacted, the more they reported feeling engaged and entertained by the experience" (Karat et. al, 2001). These results are disconcerting in that they contradict accepted learning theories that support the value of active involvement (Dewey, 1916).

A Research Study of Web Users Preferences

If people learn best in active modes but prefer passive Web experiences, how can we develop sound educational Web activities that attract and appeal to a broad audience? We decided to develop a detailed and focused pilot study of user preferences to shed more light on this complex and important issue. The primary research question was: How do people's preferences vary among types of Web-based learning activity?

We identified six activity types based on our previous Web development experience and a review of the literature (Gogg and Mott, 1993; Karat et al, 2001; Plowman, 1996b; Sumption, 2001). The six types, as described in the survey instrument, were:

- **Creative Play.** Draw a picture, write a story, make a movie, etc. Create something original based on the things you learn along the way.

- **Guided Tour.** Join an expert to explore a topic that he or she knows and loves. The guide leads you on their path through the topic.

- **Interactive Reference.** Explore a topic on your own, through informative words and pictures. Choose the links that interest you to find out what you want to know.

- **Puzzle/Interactive Mystery.** Put on your thinking cap and solve a puzzle or mystery. Put the clues together to discover the right answer.

- **Role-playing Story.** Choose your own adventure - pick a character, play a role, make decisions, and see what happens. You choose your path through the story.

- **Simulation.** Run a model of the real world and see what happens when you change things. The choices you make determine the results.

Referring back to Gammon's typology of learning, Interactive Reference and Guided Tour lend themselves primarily to cognitive learning. Creative Play, Puzzle/Mystery, Role-playing Story, and Simulation support both affective learning and developing skills. Creative Play will help learners with skills such as storytelling and art making. Puzzle/Mystery and Simulations with prediction, deduction, and other problem-solving skills. Role-playing Stories can challenge beliefs and values and help learners appreciate other people's points of view.

Methodology

In November-December 2001, we conducted a series of pilot studies with visitors to one site in order to test various versions of the exit questionnaire. The challenge was to describe the types of learning activity in such a way that preference for type of learning activity was not confounded by preference for the subject matter of the particular site or its visual appearance. We tried and eliminated screenshots of sample sites, since respondents were found to cue to content and aesthetics more than the general activity type. Long Likert scales (5- and 7- point) seemed to confuse respondents, who often indicated contradictory preferences over a series of questions.

Once we finalized the twelve-question survey, five activity sites previously developed by Educational Web Adventures, alone or in collaboration with its clients, were selected for this initial study to represent five of the six types of Web learning activity. No site exemplifying the Guided Tour was represented in the Eduweb portfolio. However, it remained in the list of types about which visitors were queried. The exit questionnaire (Appendix A) was placed on each of five educational Web sites.

In addition to the exit survey, server statistics were used to determine the duration of stay. Summary, a log analyzer (www.summary.net), generated duration charts that give a clear picture of how long users spend at each site.

A pop-up window displayed the survey on each activity site; it appeared when visitors came to the initial page, and remained behind the main browser window until the visitor clicked to leave the site. Then the survey returned to the foreground. The surveys were posted on the activity sites for 10-20 days, until 50 responses from each site were collected. Table 1 outlines the sites, types of activity, and sample size.

A control group consisting of 299 visitors to the Educational Web Adventures Web site filled out the first part of the questionnaire, which dealt with learning in general and did not reference a particular activity. Members of the control group did not engage in any of the Web activities selected for the exit survey. The purpose of the control group was to provide a measure of user preferences independent of a specific learning activity for comparison to questionnaires filled out at the activity sites.

	Type	Name	Host Site	Number
Activity Sites				
	Creative Play	A Brush with Wildlife	Nat. Museum of Wildlife Art	50
	Interactive Reference	Study Art	Sanford	50
	Puzzle/Mystery	Leonardo's Workshop	Sanford	50
	Role-playing Story	In Search of the Ways of Knowing Trail	Brookfield Zoo	50
	Simulation	Modeling Marine Ecosystems (subscription site)	JASON Project	50
Control Site				
	Control	Web Adventure directory page	Educational Web Adventures	299

Table 1. Experimental Design

In the following discussion, Control group results are compared to the Treatment population. The Treatment group consists of visitors to the five different Activity Sites. It is important to note that Treatment and Control are not used in the conventional way. We are not looking for post-treatment learning effects. Rather, we are comparing the preferences of users who have and have not experienced a particular Web activity. The purpose of the comparison is to be certain that observed user preferences are not *solely* a function of the activity in which they have just engaged.

Results

The Activity sites had significantly more children than the Control site (X^2 p = .0001). An unexpectedly large number of adult females visited the control site.

	Percents	
Demographics	Activity Sites	Control
Adults	31	49
Adult Males	11	12
Adult Females	21	37
Children	69	51
Boys	26	22
Girls	42	30

Table 2. Demographic Characteristics of Treatment/Control

	Sites (Percents)					
	Creative Play	Interactive Reference	Puzzle/ Mystery	Role-playing	Simulation	Control
Adults	25	56	32	32	13	49
Males	6	8	14	20	7	12
Females	18	48	18	12	7	37
Children	75	44	68	68	87	51
Boys	35	10	16	32	37	22
Girls	41	34	52	36	50	30

Table 3. Demographic Characteristics of Visitors by Site

The most important differences (X^2 p = <.0001) are the large percent of adult females at both the Control and the Interactive Reference sites and the large percent of children at the Simulation site. Also the number of adult males at the Role-playing site was greater than expected, as was the number of girls at the Puzzle site.

The Activity sites differ significantly from the Control site (X^2, p = .03). The respondents from Activity sites chose Creative Play and Interactive Reference more than expected. The Control site chose Puzzle/Mystery and Simulation more than expected. As will be seen below, the differences are probably due to the unexpectedly high percentage of adult females at the Control site and a significantly higher percentage of children at the Activity sites.

We see immediately that, contrary to the findings of the IBM study, Guided Tour was *the least preferred* type of Web activity for both the Treatment (Activity) and Control sites. The reasons for the difference between this and the IBM study are due to differences in both sample and methodology. IBM used adult subjects, ages 21-55. The subjects, employees and interns at an IBM research facility, were recruited to evaluate selected Web sites. IBM's sample sites dealt with the subject of music. In the study described here, children are a significant segment of the population, the user group at each site is voluntary and self-selected, and subject matter varies from site to site.

With the exception of the low scores for Guided Tour, the other types of learning activity seem to be about equal in user preference at the Control site. However, significant differences emerge when the user group is subdivided by generation and gender.

BEST COPY AVAILABLE

	Percents		Percents			
Program Type	Adults	Children	Adult Males	Adult Females	Boys	Girls
Creative Play	9	27	10	9	41	15
Guided Tour	10	7	10	10	4	10
Interactive Reference	26	11	20	29	10	13
Puzzle/ Mystery	21	18	30	19	12	24
Role-playing Story	13	23	10	14	16	28
Simulation	21	14	20	20	18	11
Totals	100	100	100	100	100	100

Table 5. Favorite Learning Activity by Age and Gender: Control Site

There are a number of differences between generations (X^2 p = .0001). Adults preferred Interactive Reference and Simulation and children preferred Creative Play and Role-playing.

Gender differences between adults were not significant. However, differences between boys and girls were significant (X^2 p = .01). Boys showed a preference for Creative Play whereas girls favored Role-playing Story and Puzzle/Mystery.

	Percents		Percents			
Program Type	Adults	Children	Adult Males	Adult Females	Boys	Girls
Creative Play	9	31	10	8	33	31
Guided Tour	15	5	15	15	4	6
Interactive Reference	40	14	25	46	11	16
Puzzle/Mystery	13	17	20	10	14	19
Role-playing Story	16	23	25	13	29	19
Simulation	7	10	5	8	10	10
Totals	100	100	100	100	100	100

Table 6. Favorite Learning Activity by Age and Gender: Activity Sites

Differences in program type preferences at the Activity sites are similar to those found at the Control site. Again there are significant differences between generations (X^2 p = <.0001). As at the Control site, adults prefer Interactive Reference and children prefer Creative Play and Role-playing. At the Activity sites, gender differences were not statistically significant.

	Sites (Percents)					
Program Type	Creative Play	Interactive Reference	Puzzle/ Mystery	Role-playing	Simu-lation	Control
Creative Play	45	10	28	16	23	18
Guided Tour	10	6	3	14	6	9
Interactive Reference	13	65	15	4	12	18
Puzzle/Mystery	10	8	21	25	12	19
Role-playing Story	18	6	18	35	32	18
Simulation	5	4	15	6	15	18
Totals	100	100	100	100	100	100

Table 7. Favorite Learning Activity by Respondent's Site

There are significant differences in program type preferences among the Activity sites (X^2 p = <.0001). As mentioned above, with the exception of a lack of interest in Guided Tour, respondents at the Control site were almost evenly divided in their preference for program type. The Activity site respondents, on the other hand, tended to prefer the type of program they were using. This is notably the case with Creative Play, Interactive Reference and Role-playing Story.

It is important to remember that the Control site data was collected as a check on the tendency of Activity site users to prefer the type of program they are using. However, the similarity in preferences between Control and Activity sites (when subdivided by generation and gender) suggests that there are patterns in preferences that transcend the particular site.

	Percents	
Locations	Activity Sites	Control
Friend's House	2	1
Home	40	40
Library	1	2
School	48	42
Work	7	11
Other*	2	4
Totals	100	100

Table 8. Location of Respondents

**Home school (3), Internet café (2) Work at Home (1), Grandparent's house (1), No answer (9). Numbers in parentheses () indicate number of responses.*

There were no significant differences between the Activity sites and the Control site in terms of where the respondents were when using the computer. Home and School were the most frequent location for both Activity and Control.

The remainder of the questionnaire dealt with users' responses to the Web activity in which they had engaged; consequently, these questions were not asked at the Control site. The tables below show results from the five Activity sites.

	Percents				
	Creative Play	Interactive Reference	Puzzle/ Mystery	Role-playing	Simulation
Friend's House	2	2	2	2	29
Home	13	76	59	40	10
Library	0	0	2	0	2
School	73	12	35	44	79
Work	13	6	2	12	2
Other*	0	4	0	2	4
Totals	100	100	100	100	100

Table 9. Location of Respondents by Site

BEST COPY AVAILABLE

The differences are significant (X^2 p = <.0001). Interactive Reference and Puzzle/Mystery are accessed most from home while Creative Play and Simulation are often school activities.

	Percents				
	Creative Play (Art)	Interactive Reference (Art)	Puzzle/ Mystery (History)	Role-playing (Natural History)	Simulation (Ecology)
Little	40	18	48	41	28
Medium	38	43	24	45	50
Lots	23	39	28	14	22
Totals	100	100	100	100	100

Table 10. Previous Knowledge of Topic

More people than expected indicated they knew "lots" for the Interactive Reference site and more indicated "little" for the Puzzle/Mystery site (X^2 p = .01).

	Percents				
	Creative Play	Interactive Reference	Puzzle/ Mystery	Role-playing	Simulation
Assigned by teacher	69	16	26	28	77
Personal interest	6	20	28	38	5
Professional interest	10	8	6	6	2
Recommended by friend	2	0	2	2	0
Recommended by parent	0	0	4	4	2
To use in a lesson	10	36	17	9	5
Other*	2	20	17	13	9
Totals	100	100	100	100	100

Table 11. Why Doing This Web Activity?

* *Creative Play: to show son (1).*
* *Interactive Reference: for a project (4), info about artist (1), daughter's homework (1), find a picture (1).*
* *Puzzle/Mystery: fun (1), no answer (6).*
* *Role-playing Story: fun (3), exploring the Web (2) recommended by teacher (1).*
* *Simulation: no answer (5).*

Teacher assignments are responsible for a high percentage of Creative Play and Simulation usage (X^2 p = < .0001). This corresponds with Table 9 that shows that these activity sites are most often used from school. Personal interest was unexpectedly high for Role-playing and the Interactive Reference site was most often accessed to use in a lesson.

	Percents
Decreased	12
Remained the same	46
Increased	43
Total	100

Table 12. Has your Enthusiasm for this Topic Changed?

Most people reported that their interest in the subject matter had stayed the same or increased after the Web-based learning activity. Results were not significantly different from site to site.

	Percents				
	Creative Play	Interactive Reference	Puzzle/ Mystery	Role-playing	Simulation
Finished it	60	54	48	75	45
Got bored	9	4	21	9	23
Got confused	2	0	5	2	18
Ran out of Time	24	20	18	9	10
Other*	4	22	9	6	13
Totals	100	100	100	100	100

Table 13. Why Respondents Left the Site

* *Creative Play: don't have Java (1), no answer (1).*
* *Interactive Reference: got what was wanted(4), didn't get what was wanted (3), Knew everything (1), going to play art game (1), no answer (3).*
* *Puzzle/Mystery: couldn't find answer (1), done with it (1), just got started (1), text too big (1).*
* *Role-playing Story: slow (1), not what I was looking for (1) finishing school work (1).*
* *Simulation: no answer (6).*

The differences here are significant (X^2 p = .007). More respondents than expected left the Simulation site because they were confused. Respondents tended to finish the Role-playing Story more often than expected and respondents left the Interactive Reference site for "other" reasons. The reasons were likely to be that the information sought was found or not found. Actually "found what was wanted" might also be interpreted as "finished", which if coded that way would increase the percent in Interactive Reference who felt they had "finished it".

The two sites with the highest percentage of users who "finished" the experience are Creative Play and Role-playing story. These are also the sites that are most preferred by children. Perhaps the goal-based structure with a clear ending is part of the appeal of the two sites.

	Percents
Worse	15
Equal	39
Better	45
Total	100

Table 14. How does this Activity Compare to Others of Its Type?

Most people felt that the Activity site was as good or better than other similar sites. There were no significant differences from site to site in users' comparison of the quality of the site.

BEST COPY AVAILABLE

Program Type	Ideal Time	Mean Time	Mean Time 2* (minus outliers)	Percent of ideal (Mean 2/Ideal)
Creative Play	21 min	6 min 19 sec	7 min 48 sec	37
Interactive Reference		7 min 25 sec	7 min 54 sec	--
Puzzle/Mystery	20 min	5 min 29 sec	7 min 42 sec	38
Role-playing Story	19 min	7 min 46 sec	12 min 24 sec	65
Simulation	25 min	16 min 40 sec	14 min 42 sec	59**

*Outliers (under 8 seconds or an hour or more) are eliminated.

** This activity required written journal entries.

Table 15. Time at Site

Additional data was gathered in the form of server logs, which were analyzed for duration information

We made a crude measure of the "Ideal Time" it takes a person to go through the activity by timing an adult clicking on and reading all the available content. We electronically collected actual time spent at the site. The mean is given first for the whole group and then minus the outliers, excluding those who left immediately (presumably because the site did not offer what they were looking for or did not have the required technology) and those who stayed more than an hour (possibly because they left their browser on the site after finishing the activity).

A comparison of mean time to Ideal Time shows that unusually long times were spent at the Role-playing Story and Simulation. However, many children were engaged in the Simulation Activity because their teacher assigned it to them (see Table

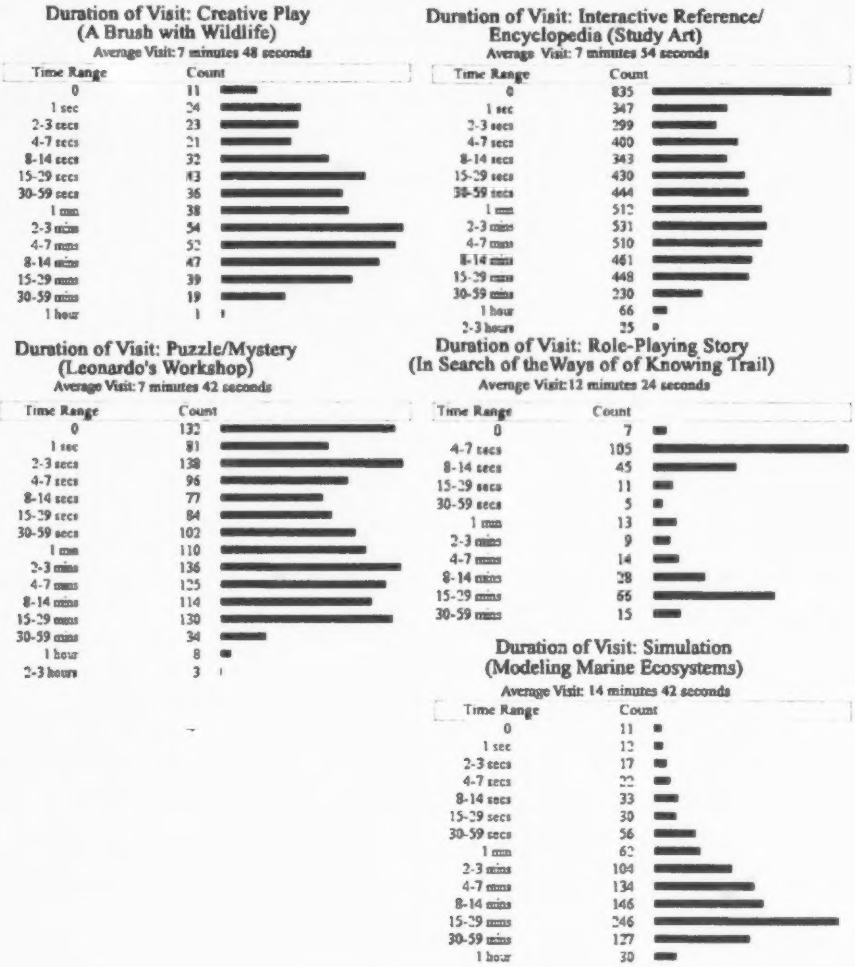

Table 16: Duration of Visit

171

BEST COPY AVAILABLE

11); they were required to complete the Simulation and submit journal entries with their conclusions.

Apart from the Role-playing Story and Simulation, the mean times are very similar, indicating that these activity types have comparable holding pow The duration charts that follow reveal in more detail the differences in holding power of the various activity types.

At most sites, with the exception of Creative Play and Simulation, a sizeable number of visitors leave within the first seven seconds. The explanation for this finding may lie in the specifics of the individual sites rather than in their Activity type.

The Marine Ecosystems Simulation was assigned as schoolwork for three-quarters of the users of the site. The Creative Play activity starts with a series of animations that may hook visitors more effectively than the text introductions to the Interactive Reference and Puzzle/Mystery sites. The Role-playing site begins with a splash/Flash plug-in detection page that turns away a sizeable percentage of visitors apparently due to the wait involved, rather than lack of the Flash plug-in (only 16% of visitors did not have the plug-in).

The remainder of each chart is more revealing. Interactive Reference shows a bell curve. The other sites have skewed distributions, indicating greater holding power, after the initial drop-off. Most striking is the curve for Role-playing Story. Half of the visitors who got past the splash page (and 25% of all visitors) stayed for at least fifteen minutes, approaching and even surpassing the ideal time (19 minutes) for that site. The curves for Creative Play and Puzzle/Mystery are less dramatic, but suggest a similar pattern.

Conclusions

Comparisons of users' responses to an exit questionnaire posted on five Web-based learning activity sites and at a control site answer some fundamental questions about users preferences for different types of Web-based learning. An earlier study of this subject conducted by IBM found that adult users prefer Guided Tours or non-interactive Web experiences. In the current study, using self-selected subjects and including a large proportion of children, Guided Tour was the *least preferred* type of Web activity for both the Treatment (Activity) and Control sites.

With the exception of the low scores for Guided Tour, the other types of learning activity seem, at first, to be about equal in user preference at the Control site. However, when the user group is subdivided by generation and gender, significant differences emerge having to do with the user's age, location and purpose for engaging in the activity.

There are significant age differences in preferences. At both the Activity and Control sites, adults prefer Interactive Reference while children prefer Creative Play and Role-playing Stories. Adult females are over-represented at the Control site. Judging by the Activity sites, many of these adult females may be teachers who use Web sites for their lessons. The Creative Play and Simulation sites used here are often school activities and are usually assigned to the children by a teacher.

Activity site respondents tended to prefer the type of program they were using. The Control site data was collected to counter this tendency. The similarity in preferences between Control and Activity sites (apart from differences in demographic composition) suggests that the observed preference patterns transcend particular sites.

Most people reported that their interest in the subject matter had stayed the same or increased after the Web-based learning activity. Also, most users felt that the Activity site was as good as or better than other similar sites. This suggests that production values did not heavily influence preference for activity type.

Implications of the Web Learning Preference Study

There are clear differences in the type of Web-based learning activity that adults prefer in comparison to children. Adults prefer the information-based activities of Interactive Reference and Simulation, whereas children, not surprisingly, are more inclined to prefer the exploratory experiences of Role-playing Story and Creative Play. The adult sites yield more straight-forward cognitive information while the sites preferred by children have strong affective components and allow more personal choice and interaction, but can lead to "dead ends" or less utilitarian solutions. Apparently, adults bring an intrinsic motivation to the learning experience. They know what they want to learn and they want to learn it in the most direct way. Children, on the other hand, need to be motivated. They respond positively to the opportunity for interaction and choice within a goal-based environment that offers them an extrinsic purpose.

Goal-based environments are advocated by Roger Shank, director of the Institute for the Learning Sciences at Northwestern University. He describes them as "Goal-Based Scenarios" (GBS)—structured learning programs that can be successful in both physical and virtual environments. The goals in these scenarios are not arbitrary extrinsic motivations, such as a good test score, prize, or reward. Rather, they stem from the activity itself—solve a crime, reach a destination, create an original artwork—which reinforces the cognitive goals of the activity. Thus, GBSs "provide motivation, a sense of accomplishment, a support system, and a focus on skills rather than facts" (Schank, 1992). In this way, they meet the basic criteria for learning of Dewey, Jonassen, and Healy cited above. They create an environment for doing and thinking, and provide both a challenge and a payoff. If designed properly, they can also connect with pre-existing knowledge and help forge an emotional connection with the subject matter.

Of the six types of Web-based learning activities explored here, two (Puzzle/Mystery and Role-playing Story) are naturally suited to the GBS approach; they inherently provide a motivation to reach a solution. Creative Play and Simulation, on the other hand, may or may not establish a clear goal. If designed as a GBS, each of these four activity types offers a goal or challenge, a payoff, structure and guidance, and some degree of interactivity. Young or novice learners who are unfamiliar with a particular learning domain need such guidance and structure to attract and hold their attention. It is interesting to note that in the Creative Play, Role-playing Story and Puzzle/Mystery activities, a plurality of users indicated they knew "little" about the subject to begin with (Table 10).

In contrast, Guided Tour and Interactive Reference are not goal-based scenarios. There is no payoff or achievement for completing the activity; learners must bring their own intrinsic motivation to the task. Users of these Web sites may have more expertise in the subject as well. On the Interactive Reference site, more people than expected indicated they knew "lots" about the subject (Table 10).

Within the structure and guidance provided by GBSs, young learners prefer some degree of freedom. Creative Play and Role-playing Stories, both preferred by children, offer a series of choices in the path of the activity and some control over the outcome. Puzzle/Mystery, which was less favored, and Guided Tour, the least popular type, offer only one outcome and less opportunity for personal involvement.

This brings us to another important dimension: whether the outcome is determined (created by the site developers) or user-created. Different pedagogies underlie the two types of outcome. Discovery learning lends itself to puzzles and mysteries, with their single correct solution (determined), while constructivism supports user-created outcomes that allow more personal choice and involvement. Most of the activities rely on determined outcomes, but Creative Play, and to a lesser extent Simulation, permit learners to create their own outcomes, be it a picture, a story, or a unique configuration of the variables in the simulation.

Putting the Results into Practice

First, developers of educational Web activities must decide whether their primary audience is adults or children, since the two groups have different learning preferences. Adults prefer reference sites. Children prefer goal-based scenarios, particularly Creative Play and Role-playing Stories. These activity types offer an appealing middle ground between the highly structured and constrained Guided Tour approach, and the "explore according to your own interests" reference format typical of much interactive multimedia. If a site must appeal to both children and adults, developers should consider a dual approach, combining reference and play (as in Leonardo's Workshop and Study Art).

Second, developers must decide on a pedagogical approach. Guided Tour and Interactive Reference sites typically provide a traditional expository/didactic approach to learning. Puzzle/Mystery activities reflect a discovery learning orientation, in which a single correct solution or conclusion is the goal. Creative Play takes a constructivist approach, encouraging open-ended experience. Simulations and Role-playing Stories can take a variety of forms, from discovery to constructivism.

Developers should also consider their audience's expertise in the subject. Expert learners with existing interest in the domain are more likely to favor interactive reference sites. Novice learners, regardless of age, are more likely to need and prefer a guided experience to introduce them to the subject and motivate them to learn more about it.

The audience's degree of expertise also affects the learning goals of a Web activity. For novices, affective learning experiences can inspire interest and curiosity in a subject, while skill-building activities help them to develop the ability to pursue further understanding. For experts who already have the knowledge base and skills to tackle the subject, interactive reference sites can provide a satisfying cognitive learning experience. For novices, goal-based scenarios can combine affective and skill building learning experiences with cognitive learning.

Of course, just because a Web activity attracts and holds users' interest doesn't mean it is achieving its educational goals. Evaluating learning outcomes is no easy task, and a matter for another study. However, a Web activity or any other learning activity must first attract and hold the interest of learners in order to have the opportunity to achieve its learning objectives.

References

Crawford, C. (1982). *The art of computer game design*. Reprinted online at:

http://www.vancouver.wsu.edu/fac/peabody/game-book/Coverpage.html

Dewey, J. (1916). *Democracy and education: An introduction to the philosophy of education.*

Gammon, Ben. (2001). *Assessing learning in museum environments: A practical guide for museum evaluators*. Unpublished Science Museum (London) report.

Gogg, T. J., Mott, J. R. A. (1993). Introduction to Simulation. *Proceedings of the 1993 Simulation Conference*, ed. G.W. Evans, M. Mollaghasemi, E. C. Russell, W. E. Biles.

Healy, Jane M. (1994). *Your Child's Growing Mind: A practical guide to brain development and learning from birth to adolescence*. New York: Doubleday.

Hein, G. (1998). *Learning in the museum*. London: Routledge.

Jonassen, David H., Kyle L. Peck, Brent G. Wilson. (1999). *Learning with technology: A constructivist approach*. Upper Saddle River: Merrill.

Karat, C., Karat, J., Claudio Pinhanez, Mark Podlaseck, John Vergo, Doug Riecken, Renee Arora, Thomas Cofino. (2001). Less clicking, more watching: Results from the user-centered design of a multi-institutional web site for art and culture. *Museums and the Web 2001: Selected Papers from an International Conference*, ed. D. Bearman, J. Trant. Pittsburgh: Archives & Museum Informatics.

Karat, C., Claudio Pinhanez, John Karat, Renee Arora, & John Vergo. (2001b). *Less clicking, more watching: results of the iterative design and evaluation of entertaining web experiences*. Interact '01.

Murray, J. H. (1997). *Hamlet on the holodeck: The future of narrative in cyberspace*. Cambridge: MIT Press.

Laurillard, D.M. (1996). *Multimedia and the learner's experience of narrative*. Keynote speech for Ed-Media, Boston 1996. http://meno.open.ac.uk/homerpub.html.

Loving, C. C. (1997). From the summit of truth to its slippery slopes: Science education's journey through positivist-postmodern territory. *American Educational Research Journal* 34(3): 421-452.

National Research Council. (2000). *Inquiry and the national science education standards: A guide for teaching and learning*. Washington: National Academy Press.

Plowman, L. (1996). Designing interactive media for schools: A review based on contextual observation. *Information Design Journal* 8 (3): 258-266. http://meno.open.ac.uk/meno-rev-idj.html.

Plowman, L. (1996b). Narrative, linearity and interactivity: making sense of interactive multimedia. *British Journal of Educational Technology* 27 (2): 92-105.

Prensky, M. (2001). *Digital game-based learning*. New York: McGraw Hill.

Schaller, David T., and Steven Allison-Bunnell (2001). *Developing Goal-Based Scenarios for Web Education*. http://www.eduweb.com/goalbasedscenarios.html

Schank, R. 1(992). *Goal-based scenarios*. Chicago: Northwestern University Institute for the Learning Sciences. http://cogprints.soton.ac.uk/documents/disk0/00/00/06/24/cog00000624-00/V11ANSEK.html.

Spectre, M. and M. Prensky. (ND). *Theoretical underpinnings of games2train's approach*. http://www.games2train.com/site/html/theory.html.

Sumption, K. (2001). Beyond museum walls — A critical analysis of emerging approaches to museum web-based education. *Museums & the Web 2001 Proceedings*, ed. D. Bearman, J. Trant. Pittsburgh: Archives & Museum Informatics. http://www.archimuse.com/mw2001/papers/sumption/sumption.html.

Tipping, P.D. & Graesser, A.C. (1996). *Active learning, passive learning, software, and texts: Does it really make a difference?* Paper presented at the Sixth Annual Meetings of the Society for Text and Discourse, San Diego, CA.

Turkle, S. (1997). Seeing through computers: Education in a culture of Simulation. American Prospect. 8 (31). http://www.prospect.org/print/V8/31/turkle-s.html.

Veenema, S. and H. Gardner. (1996). Multimedia and multiple intelligences. *American Prospect*, 7(29). http://www.prospect.org/print/V7/29/veenema-s.html.

Weatherford, J. ND. *Instructional Simulations: an overview*. (ND). San Diego State University Dept. of Educational Technology. http://coe.sdsu.edu/eet/Articles/instrucsimu/start.htm

Appendix A: Web Activity Preference Survey

When you are done exploring the *Study Art* section of this Web site, we would like to ask you a few questions. Your responses will help us develop better Web sites in the future. Your responses are anonymous and will be kept confidential.

Please complete this survey when you are **leaving *Study Art*.**

Are you:
Male Female

How old are you?
5-8 9-10 11-13 14-18 19-22 23-35
36-49 50+

Where are you now?
Home School Work Library Friend's house
or Other: [enter text]

How much did you know about art before you came to this Web site?
Very littleMedium A lot

Why were you looking at this Web site?
Assigned by teacher
Recommended by a friend
Recommended by parent
To use in a lesson
Personal interest
Professional interest
or Other: [enter text]

Why are you leaving this Web site?
Found what I was looking for Ran out of time
Got bored Got confused
or Other: [enter text]

Study Art is an Interactive Reference/Encyclopedia Web site. How does it compare to other such sites you've seen on the Web?
Much worse Equal to Much better

What would you say Study Art is about?
[enter text]

What is one thing you learned that you didn't know before?
[enter text]

Now that you're leaving Study Art, has your enthusiasm for art changed?
Decreased Stayed the same Increased

What is your favorite type of computer learning activity? There are six types of activities below. Please choose one:

Guided tour. Join an expert to explore a topic that he or she knows and loves. The guide leads you on the path they chose through the topic.

Interactive reference/encyclopedia. Explore a topic on your own, through informative words and pictures. Choose the links that interest you to find out what you want to know.

Role-playing story. Choose your own adventure—pick a character, play a role, make decisions, and see what happens. You choose your path through the story.

Creative play. Draw a picture, write a story, make a movie, etc. Create something original based on the things you learn along the way.

Simulations. Run a model of the real world and see what happens when you change things. The choices you make determine the results.

Puzzle or interactive mystery. Put on your thinking cap and solve a puzzle or mystery. Put the clues together to discover the right answer.

If you have any comments about this survey or Study Art, please write them here:
[enter text]

Evaluating The Features of Museum Web Sites (The Bologna Report)

Nicoletta Di Blas, HOC-DEI, Politecnico di Milano, Maria Pia Guermandi, IBC, Istituto Beni Culturali, Emilia Romagna, Carolina Orsini, Università di Bologna, Paolo Paolini, Politecnico di Milano, Italy

Abstract

MiLE (Milano – Lugano Evaluation Method) is an innovative method for evaluating the quality and usability of hypermedia applications. This paper focuses upon the specific "module" of MiLE concerning cultural heritage applications, synthesizing the results of research carried on by a group of seven museum experts of Bologna (Italy), with the joint coordination of IBC (Institute for the Cultural Heritage of the Emilia Romagna Region) and Politecnico di Milano. The "Bologna group" is composed of different professional figures working in the museum domain: museum curators of artistic, archaeological and historical heritage; museum communication experts; Web sites of cultural institutions' communication experts. After illustrating the general features of MiLE and the specific features for Cultural Heritage, we will briefly show a few of the results which are to be published in the "Bologna Report".

Keywords: usability, inspection method, cultural heritage, users' scenarios

1. MiLE in a nutshell

MiLE is based upon a combination of Inspection (i.e. an expert evaluator, systematically exploring the application) and Empirical Testing (i.e. a panel of end users actually using the application, under the guidance and the observation of usability experts). If this combination of the two methods is not new (several usability methods propose, in fact, a similar combination), the innovation of MiLE comes from the set of guidelines being used for making both inspection and empirical testing more effective and reliable. In extreme synthesis, we introduce two specifics (heuristic concepts):

Abstract Tasks, ATs in short, used for inspection. They are a list of generic actions (generic in that they can be applied to a wide range of applications) capable of leading the inspector through the maze of the different parts and levels an application is made of, as the Ariadne's thread. MiLE, in fact, provides inspectors with some guidelines that draw their attention to the most relevant features of the application.

Concrete Tasks, CTs in short. They are a list of specific actions (specific in that they are defined for a single application) which users are required to perform while exploring the application for the empirical testing.

Inspection is the focus of this paper, and we will not further explore the issues concerning empirical testing. One contribution of MiLE is the emphasis on the need for separating different levels of analysis: technology, navigation, content, illocutionary force, graphic, etc. For each level a library of Abstract Tasks has to be prepared, when building the method, in order to support the inspection. The Abstract Tasks are nothing but the marrow of each level's experts' knowledge. For some levels (e.g. graphic or navigation), the abstract tasks can be generally independent from the specific application domain; for other levels (e.g. content) we have different tasks according to the application domain (i.e., specific tasks for the cultural heritage domain, for the e-commerce domain, and so on).

The inspector has to understand the client's communicative goals, combine them with the intended users' probable requirements, and then select the appropriate set of tasks to perform. If, for example,

BEST COPY AVAILABLE

we have to evaluate a museum's site that is specifically meant to attract visitors to the real museum, we – as inspectors –concentrate on all those tasks that at the content level involve the practical services parts of the site (opening hours, "buy a Ticket", etc.).

When performing inspection, the inspector has to check a list of attributes concerning the different facets of usability/quality (e.g. richness, completeness, etc.). For each attribute (in relation to a specific AT), a score must be given. After the scoring phase is over, the set of collected scores is analyzed through "weights" which define the relevance of each attribute for a specific goal (or, technically speaking, for a "user scenario").

Weighting allows us a clean separation between the "scoring phase" (using the application, performing the tasks, and examining them) from the "evaluation phase" in a strict sense, where different possible usages are considered. Let us introduce a simple example: assume that a navigation feature (e.g. using indexes) is not very powerful, but very easy to learn. What should the evaluation be? With MiLE the inspector can provide a score (e.g. 9/10 for "predictability" and 2/10 for "powerfulness") for the navigation. Later, figuring out two different user scenarios (e.g. casual users and professional users), the evaluator (possibly different from the inspector), can assign two different pairs of weights to the attributes "predictability" and "powerfulness". The weights, for example could be <0.8 (predictability), 0.2 (powerfulness)>, for casual users, or <0.1 (predictability), 0.9 (powerfulness)> for professional users. The weighted score for the navigation feature is very different of course (7.6 for casual users and 2.7 respectively), but it reflects the different users' scenarios. The inspector could therefore conclude that the application (at least for this feature) is well suited for casual users, while it is somehow ineffective for professional users. Trying different weighting systems allows the evaluator to test different user scenarios using the same set of scores derived from the inspection.

In short, an inspection with MiLE requires the following steps:

- selection of the relevant portion of the application (based upon considerations that we can't investigate here);
- selection of the Abstract Tasks that are relevant for the intended user scenarios;
- execution of the Abstract Tasks, providing scores for each attribute;
- for each user scenario
- weighting of the attributes and of the tasks chosen (a certain task can be more relevant than another);
- production of quantitative evaluation measures (applying weights to scores)

The reliability of the method has proved to be very high. Execution of the Abstract Tasks (at navigation and content level) allows producing more reliable evaluation results and helps spot unexpected usability problems (inconsistencies, lack of clarity, etc.). Even "at-first-sight agreeable" sites, when put to the test through a systematic inspection "à la MiLE", may reveal weaknesses and defects.

Inspection already provides valuable evaluations; in some cases, however, panels of users may be required for double checking. When empirical testing is required, users are given a list of concrete tasks, i.e. a list of specific actions that they are asked to perform. Concrete tasks definition (different for each case) is based upon the results of the inspection, which has identified portions of the application, tasks and attributes that need special attention. "Real users" can fine-tune the inspector's observations, confirming them (very often), or dissenting from them (seldom) or spotting additional problems. We now discuss the above outlined approach applied to cultural heritage applications.

2. Guidelines for Evaluating Cultural Heritage Applications

Some of the features (such as navigation or layout) of an application can be examined largely independently from a specific application domain; other features, such as content or functions offered to the users, require a different evaluation schema for each application domain. In order to explore functions and contents for museum Web sites (a specific subdomain, within the larger domain of cultural heritage applications), a specific panel of "experts" (the so-called "Bologna group") has been created, with a partnership between Politecnico di Milano and "Istituto Beni Culturali", a regional organization supervising cultural heritage activities in a large region, Emilia Romagna, with headquarters in Bologna. The group is composed of museum curators (both of archaeology, modern and contemporary art museums and galleries), museum communication experts, and researchers of new technologies for cultural heritage.

The first step of the Bologna group was to identify the main pieces of a generic Museum Web site. In order to avoid the danger of "wish listing" the sum of what everybody could foresee as the "ideal Web site", we took an empirical standing: we selected a large number of sites and considered them to be the "universe of discourse". The resulting model is therefore a synthesis of contents and features found in those sites. At this stage of the research, we have listed more than a hundred "elementary" constituents, organized into three main groups:

A. site presentation: general information about the Web site;
B. museum presentation: contents and functions referring to a "physical museum" (like "arrows" pointing to the real world);
C. the virtual museum: contents and functions exploiting the communicative strength of the medium.

A further analysis has allowed us to detect "high level" constituents such as, *collections, services, promotion,* which gather the elementary constituents (a full account of all the pieces of the model can be found as an Appendix to the electronic edition of this paper).

The next job has been to define a set of users' scenarios as a way to build a library of suitable ATs. A user scenario, in this context, is a pair <user profile, operation (that users may wish to perform)>. In simple words, we tried to identify a number of user profiles (culture, expertise, interests, etc.) and for each of them we tried to provide a set of meaningful answers to this simple question: "What might a user (with this profile) want to do with the application?"

Therefore the tasks are coupled to user profiles, in the sense that a given task may be interesting for a given profile, and/or meaningless (or irrelevant) for a different profile. When the inspectors perform an inspection, they will learn from their customers who the intended users of the application are and will concentrate on those tasks likely to be performed by these customers. In any case the Inspectors will be free to create new tasks that fit better the communicative goals of the application, as long as they follow the guidelines of the method and its "philosophy".

Overall we have classified ATs according to two different dimensions: "scope" and "concern". Possible values for scope are the following: Narrow (a specific item is interesting), Complex (several items are interesting), and General (a generic overview). Possible values for concern are Practical Info (the user wants to gather useful information), Operational (the user wants to do something) and Cognitive (the user wishes to learn something).

Table 1 shows some examples of AT, classified accordingly:

Regarding the users, we took into consideration a number of variables, such as age, expertise, professional interest (e.g. school students, fine arts students, fine arts experts, tourists, etc.) Each relevant user profile is based upon a number of these variables.

On-going research work consists of identifying the largest possible number of ATs. More than fifty ATs have been identified so far (while 49 ATs where

	NARROW **Specific item**	**COMPLEX** **Several items**	**GENERAL** **General outlook**
Practical Info	search the hours of opening of an exhibition in a specific date	organize a visit to the museum	get an idea of the overall organization of the museum
Operational	buy a ticket	reserve a group visit with a museum guide	get of an idea of all the possible arrangements for group visits
Cognitive	find a specific work of art and its description	find all the works of the "Florentine school"	get an overviews of the collections of the museum

Table 1: Some examples of AT

identified for navigation, in another research), and more than double that number will be the likely result.

Regarding the list of attributes to be scored during the inspection, we started with the idea that they would be different for each AT. At the moment, however, we have developed the following list which seems to be applicable (with minor problems only) virtually to any AT:

- Efficiency: the action can be performed successfully and quickly
- Authority: the author is competent in relation to the subject
- Currency: the time scope of the content's validity is clearly stated. The info is updated.
- Consistency: similar pieces of information are dealt with in similar fashions
- Structure effectiveness: the organization of the content pieces is not disorienting

- Accessibility: the information is easily and intuitively accessible
- Completeness: the user can find all the information required
- Richness: the information required is rich (many examples, data…)
- Clarity: the information is easy to understand
- Conciseness: the basic pieces of information are given; texts are not too long and redundant
- Multimediality: different media are used to convey the information
- Multilinguisticity: the information is given in more than one language

3. Some Examples

In this section we will introduce a few examples of inspection to help the reader grasp how our method works. The examples are very simple, and are taken from actual Web sites. We hope that in the period between the writing of this paper and the reading of it by a user, the Web sites will not be modified, so that the readers may try directly to "inspect" them. (The impossibility of "freezing" Web sites, in practice, makes it difficult to develop examples of inspection that could maintain their validity over a long span of time.)

Example 1 (practical Info AT)

find the events/exhibitions/lectures occurring on a specific date in a real museum

The user's scenario for this task is that of well-educated French-speaking tourists (who can speak English too), first-time visitors to the site, who know that on March 9th (Saturday), 2002, they will be in the town where the real museum is actually located. Therefore they would like to know what special exhibitions or activities of any kind (lectures, guided tours, concerts) will take place in that day.

We performed this task on many different Web sites, and we describe here our findings for the Louvre site (www.louvre.fr) and the Royal Ontario Museum site (www.rom.on.ca), on the basis of an inspection that took place on February 13th, 2002. The focus of our attention is the section named "information about museum activities and events" in our schema.

The relevant attributes that we use for this brief example are the following:

(A1) currency of the information;

(A2) quality of the organization of the information, since users are looking for operational support;

(A3) multilinguisticity, fundamental for an international audience;

(A4) richness of the information provided, very important in order to make understandable the potential interest of the events.

The Louvre Web site offers a choice among four languages (French, English, Spanish, Japanese). In the home page we find (on the left menu) three relevant links: "Expositions", "Auditorium" (a rather "obscure" name, for it refers to a specific place in the real museum, but its actual meaning can be understood only by second-time users) and "Visites – conférences et ateliers".

If we click on "Expositions", we get a list of the exhibitions currently available or coming soon; "temporal windows" allow users to easily select the exhibition fitting their needs.

If we click on "Auditorium", we find a wide choice of activities: « agenda », « concerts », « cinéma muet en concert », « classique en images », « colloques », « conférences », « les enfants au Louvre », « films », « lectures », « midis du Louvre » (conferences and activities taking place at midday), « musée-musées », « musique filmée », « l'oeuvre en direct ».

From a graphic point of view, the first element of the list ("agenda") looks exactly the same as all the others; only after a short exploration do the visitors discover that under the voice "agenda" they can find, arranged in chronological order, all the pieces of information that are available under the other twelve voices.

If eventually we click on "Visites – conférences et ateliers" we have to choose whether we're interested in "visites-conferences" or "ateliers" and whether we're adults or children. The "visites-conferences" are furtherly divided into these categories: « visite découverte », « visite d'une collection », « visites thématiques », « thèmes du lundi soir », « monographies d'artistes », « une heure-une œuvre », « visite d'exposition temporaire », « promenades architecturales ». Clicking on "promenades architecturales" we discover that on March 9th (Saturday) there is a special guided tour called "Le Louvre, l'oeuvre et le musée", of which we find no mention in the agenda. Therefore, even once one has found out the "collective" character of the agenda-page, a brief exploration of other related links makes it clear beyond doubt that the agenda-page is not exhaustive as regards all the special events taking place at the Louvre on a specific date. On the whole, we can say that although the information is well updated and exhaustive once found, too many paths have to be trodden in order to get to the point.

The Royal Ontario Museum shows a better solution: in the home page we find "what's on calendar", a link that leads us to a page showing a big calendar of the current month, divided into cells; all the events taking place on a specific day are listed in the day's cell and are all linkable to a description's page. Previous or next months (or even years) are shown on demand by clicking on the arrows on the top of the page. The task can be performed quickly and easily. Note that if our French-speaking tourist doesn't speak English at all and decides to enter the French version of the site, he has to choose the item "expositions", leading to a much less appealing list of the present, future and past exhibitions very similar to the Louvre's. For the scoring we have considered the English version.

Table 2 below synthesizes our scoring and evaluation.

We do not ask the reader to agree with our scores (we may be poor inspectors) but to appreciate the method on a number of issues:

a) We are evaluating a specific task and not expressing a global evaluation; in addition, we are scoring each single attribute. This level of detail introduces two advantages: precision of the feedback to application designers and possibility of pinpointing the causes for possible discrepancies among different inspectors.

b) Through weights we can take into account the specific objectives for the (portion of the) application. In the example above, we gave great relevance to attributes A2 and A1, and minor relevance to A3 and A4.

c) Global concise evaluation can be obtained trough combining the evaluation for each attribute (as in the above table), and/or combining the evaluation for the different ATs (again using weights in order to attribute different relevance to each AT).

		A1	A2	A3	A4	Global score for this AT
Scores	Louvre	8	2	6	7	5.75 (just average score)
	ROM	10	10	6	8	8.50 (just average score)
Weights		0.3	0.5	0.1	0.1	
Weighted scores	Louvre	2.4	1.0	0.6	0.7	**4.70** ("weighted" average)
	ROM	3.0	5.0	0.6	0.8	**9.40** ("weighted" average)

Table 2: The scores and the evaluation for "a visit in a given day"

d) Different systems of weights can be used in order to take into account different user profiles.

Example 2 (Cognitive AT)

find all the works of an artist shown in the site

This task might be performed by a high-school student looking for some information about an artist he's currently studying at school; let's say Giovanni Battista Tiepolo. He finds out that some ofTiepolo's works are kept by the Met Museum (www.metmuseum.org) and by the Hermitage Mu"Tiepolo" (or the full name "Giovanni Battista Tiepolo", in order to avoid the mixing between Giovanni Battista's and Giovanni Domenico's works). This more precise searching tool gives as a result the list of the 23 Giovanni Battista Tiepolo's works of art shown in the Web site. For each of the works, we have also the basic data, a description and the possibility of zooming the image.

If we decide to ignore the search and to find what we're looking for by navigating the site, then we have to reach the sub-section "European paintings" of the section "the collection"; here we find a brief introduction in which there's a mention of our author and a link. Following the link we are shown a single work ofTiepolo ("Allegory of the planets and continents"), but having entered the guided tour of the department's highlights, if we click on the "next" or "previous" buttons, then we find other artists' works of art but no more Tiepolos. In order to perform our task, we have either to check one by one the 2275 items preserved by the department (clicking on "entire department") or, if we don't want duction to Western European painting with some little icons on the right side, one of them representing the painting by Tiepolo: "Maecenas Presenting the Liberal Arts to Emperor Augustus". A link leads to a bigger image and a description. As an alternative, we can decide to browse the "digital collection" by type of art work ("paintings prints and drawings") and artist, getting nine results (corresponding in this case to the nine works by Giovanni Battista Tiepolo).

On the whole, we can say that both sites permit reaching the wished information only by using the search engines: this can be considered a sign of poor organization of the information.

For the Metropolitan, the collection's search engine must be used and not the "main" search; otherwise the user may get completely lost! The Hermitage search engine doesn't distinguish between the works of Giovanni Battista and Giovanni Domenico Tiepolo. The Metropolitan Museum offers a good description of all the items, whilst the Hermitage offers only one single work's description (the others are simply "shown").

Table 3 synthesizes our scoring and evaluation.

We should note first, that in this case the weights do not change the relative evaluation of the two sites, but rather reduce both of them, given the high relevance assigned to A4. Secondly, we can see that if more details about navigation are wanted, then a different level of analysis should be entered: we have devised nearly 50 different tasks in order to inspect navigation features precisely.

4. Conclusions and Future Work

The general distinctive features introduced by MILE can be synthesized as follows:

- Efficient combination of inspection and empirical testing
- Use of Abstract Tasks, ATs, as guidelines for inspection

		A1	**A2**	**A3**	**A4**	**Global score for this AT**
Scores	MetMus.	7	7	7	3	6.00 (just average score)
	Hermitage	7	7	4	3	5.25 (just average score)
Weights		0.15	0.15	0.25	0.45	
Weighted scores	MetMus.	1.05	1.05	1.75	1.35	**5.20** ("weighted" average)
	Hermitage	1.05	1.05	1.00	1.35	**4.45** ("weighted" average)

Table 3: The scores and the evaluation for "all the works of a given artist"

- Use of Attributes as a way to detail scoring
- Use of Concrete Tasks, CTs, as guidelines for empirical testing
- Use of weights as a way to translate scores into evaluation
- Use of user profiles in order to assign weights

The specific contribution of the Bologna group (there is also another research group, coordinated by The Museum of Science and Technology of Milan, examining the same issue for scientific and technical museums) is described in this paper. Our task has been the identification of a general framework for defining a set of AT suitable for Art Museum Web sites The framework is the result of an extensive analysis of several Web sites which are now the objective of our trial inspection.

The current work consists of identifying, through the ATs, the "universe of possible functions" that a museum Web site should support; the next step will be to pair user profile features with ATs. The goal is to generate an overall schema showing what type of user is interested and in what information/action. The combination of user-profile/AT is what we mean by User Scenario; therefore, we could also say that we are trying to build a large set of possible user scenarios for museum Web sites.

We aim at providing a contribution to the community of people interested in museum Web sites (museum curators, designers, Web managers, etc.), sharing our understanding of what it means to evaluate quality and usability of "virtual artifacts".

Since the amount of work to be performed is immense, and we would like to generate a discussion in a large community, the authors encourage all interested persons to contact them in order to enlarge the scope and the validity of this research in evaluation.

5. References

Costabile, M.F., Garzotto, F., Matera, M., Paolini, P. The SUE Inspection: A Systematic and Effective Method for Usability Evaluation of Hypermedia. In IEEE *Transactions on Systems, Man, and Cybernetics* - In print.

Costabile, M.F., Garzotto, F., Matera, M., Paolini, P. Abstract Tasks and Concrete Tasks for the Evaluation of Multimedia Applications Presented at the Int. Workshop on Theoretical Foundations of Design, Use, and Evaluation, Los Angeles, CA, USA, 1998.

De Angeli, M. Costabile, F., Garzotto, F. Matera, M., Paolini, P. On the advantages of a Systematic Inspection for Evaluating Hypermedia Usability. Iin *International Journal of Human Computer Interaction*, Erlbaoum Publ. - In print.

De Angeli, M. Matera, M. F. ,Costabile, F. ,Garzotto, F., and Paolini, P.. Validating the SUE Inspection Technique. In *Proc. AVI, 2000*, pp. 143-150.

Garzotto, F. & Matera, M. A Systematic Method for Hypermedia Usability Inspection. In *The New Review of Hypermedia and Multimedia*, vol. 3, pp. 39-65, 1997.

Garzotto, F., Matera, M., Paolini, P. A Framework for Hypermedia Design and Usability Evaluation. In *Designing Effective and Usable Multimedia Systems*, P. Jonhson, A. Sutcliffe, J. Ziegler Eds., Boston, MA: Kluwer Academic, 1998, pp. 7-21.

Garzotto, F., Matera, M., Paolini, P. Abstract Tasks: a Tool for the Inspection of Web-sites and Off-line Hypermedia, in Proc. ACM HT, 1999, pp. 157-163

Acknowledgements

We wish to acknowledge the work of the other members of the Bologna group, who made this (still on going) research effort possible. We therefore warmly thank Dede Auregli (Galleria d'Arte Moderna di Bologna), Gilberta Franzoni (Musei Civici di Arte Antica di Bologna), Paola Giovetti (Museo Civico Archeologico di Bologna), Laura Minarini (Museo Civico Archeologico di Bologna), Federica Liguori (Politecnico di Milano), Uliana Zanetti (Galleria d'Arte Moderna di Bologna).

Editor's Note: the Contents Survey Schema used in this study can be found in the electronic edition.

Interpenetration

Towards Tangible Virtualities: Tangialities

Slavko Milekic, M.D., PhD, The University of the Arts, USA

Abstract

Rapid proliferation of different types of interaction devices that use more natural channels (voice, touch, gesture) for interfacing with the digital medium illustrates the trend (and need) towards the creation of more 'humane' interaction mechanisms. However, the current historical paradox is that modern technological advances are dramatically ahead of our understanding of their possible uses and meaning on a conceptual level.

In this paper I will present an overview of some of the currently available interaction technologies, the conceptual barriers that limit their use and the case for the creation of interaction mechanisms that make abstract (virtual) information more tangible.

Keywords: interfaces, interaction, haptic interfaces, perception

Introduction

The meanings associated with the adjective "tangible" in an on-line version of the Merriam-Webster dictionary (see references for URL) include:

> **1 a** : capable of being perceived especially by the sense of touch : **PALPABLE**
> **b** : substantially real : **MATERIAL**
>
> **2** : capable of being precisely identified or realized by the mind <her grief was tangible>

The listed synonym for "tangible" is PERCEPTIBLE, which in turn, has the following synonyms: SENSIBLE, PALPABLE, TANGIBLE, APPRECIABLE, PONDERABLE. The evolution of the term, starting with sense percepts related to the sense of touch and ending with precise mental identification and realization of abstract concepts (like 'grief' in the definition above), corresponds, more or less, to my view on this topic. In this paper, I would like to make a case that association of virtual and abstract information with multimodal sensory experiences creates a new layer of knowledge and action spaces that is more natural and efficient for humans. These in-between domains, where interactions with virtual data produce tangible sensations, I dubbed tangialities (see Figure 1.).

Please note that the way I define the term tangiality includes all sensory modalities and is not reduced just to those related to the sense of touch (haptic, cutaneous, tactile).

Our body may be considered as the first interface between ourselves and the real world. The interactions were guided by our goals (intentions), carried out through actions, and repeated or corrected based on perception of the consequences of actions (observations).

Body actions were soon enhanced through the use of tools (artifacts). Norman (Norman 1991) introduced the concept of a cognitive artifact, as a tool that enhances cognitive operations. Although the enhancement of body actions is sometimes achieved by sheer magnification (using a lever, or inclined plane), the enhancement of cognitive operations is most often the consequence of *changing the nature*

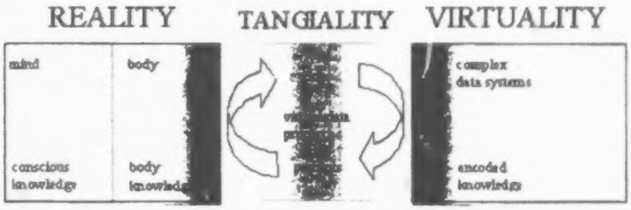

Figure 1. Representation of different action/knowledge domains

BEST COPY AVAILABLE

of the task. An example of a cognitive artifact that enhances our ability to memorize and recall events is a personal calendar. Instead of trying to rehearse and memorize all of the events for weeks to come, we have to remember only to write them down into the calendar, and to consult the calendar every day. In the context of tangialities, cognitive enhancements are also the consequence of changing the nature of the task. Most often a change is in shift from relying on formal, abstract operations as a means of gaining knowledge, to direct manipulation of data (properties) with instantaneously observable results. For example, in order to answer the problem illustrated in Figure 2., "are the dimensions of the smaller cubes exactly one-half of the larger one?" we may use conventional knowledge of algebra and solve the problem. Direct manipulation approach would be just to juxtapose two smaller cubes next to the larger one and the answer becomes self-evident. Note that the formal solution can be made harder by choosing different dimensions (for example, the height of the big cube could be 8.372914) but this does not influence the direct manipulation solution.

I would like to add another word of clarification. For the purposes of this paper, I am not going to address a very fruitful area of research often referred to as "tangible interfaces" (for example, see Ullmer, B., Ishii, H. 2000, and Patten, J., Ishii, H., Hines, J., Pangaro, G. 2001). The cornerstone of this approach is in using real objects with desired physical (manipulable) properties as data representations/containers. These objects embody computation regardless of whether they are connected to a computer or not. Although the areas of development of tangible interfaces and tangialities overlap, and will probably merge in the future, in this paper I will focus on procedures that endow data representations with tangible properties and thus make manipulations carried out on data available to our senses.

Direct Manipulation

The first human-computer interfaces were abstract, efficient and accessible only to expert users. They involved learning the vocabulary and syntax of a command language which was then used to initiate some operations on the digitally stored data, and often one needed to issue a separate command to see the results of the previous one. There was no continuity of interaction - once the command was issued, there was no way of interfering with the process (short of aborting it). There was also no sensory feedback that would provide relevant information about the operation on an experiential level.

One of the first examples of a tangiality domain was the introduction of Graphical User Interface (GUI) and the concept of direct manipulation. "Direct manipulation", a somewhat misleading term, was introduced by Ben Shneiderman in 1983 to describe what we take today to be an integral part of human-computer interaction - the use of a mouse (cursor) for pointing at and manipulating graphically represented objects. The crucial characteristics of direct manipulation are: a) continuous visibility of the manipulated object; b) all the actions carried out on the objects are rapid, incremental and reversible; and c) the consequences of actions are immediately visible (Shneiderman, 1983, 1998).

What makes direct manipulation a tangiality domain is the fact that it provided continuous sensory input (visual and kinesthetic feedback from hand-on-mouse positions) while acting on abstract parameters (like location coordinates, adjacency, parallelism) of digitally represented data. In spite of the fact that the output in direct manipulation depended on a single sense (vision), it truly revolutionized human-computer interaction.

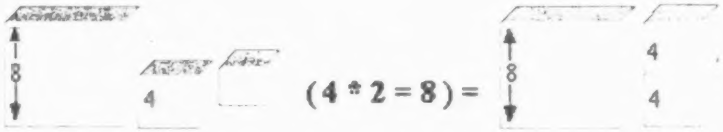

Figure 2. Getting the result using abstract operation or direct manipulation

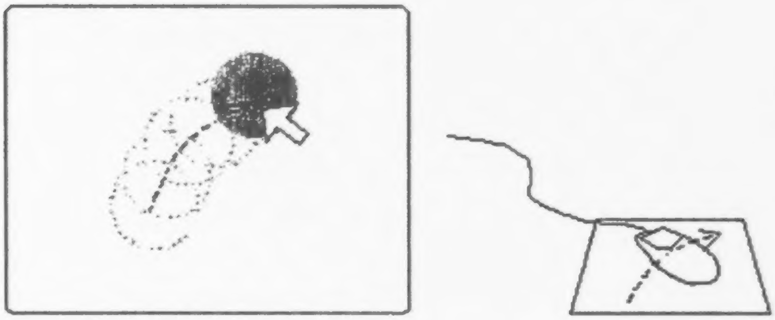

Figure 3. "Direct manipulation". Notice that in spite of the term the movements of the mouse are (indirectly) mapped to the screen coordinates.

Suddenly, anyone who could see and make hand movements could use the computer. However, it is the very success of the direct manipulation paradigm that is now one of the obstacles to creating even more efficient interfaces.

Problems with the Traditional Interface

As Malcolm McCullough aptly put it in his book "Abstracting Craft" (McCullough, 1996), one of the problems with the traditional "point-and-click" interface was the increased separation of the hand and the eye. In performing operations on digital data, the eye was given a mayor role of identifying, focusing, monitoring and interpreting, while the hand was reduced to performing simple repetitive gestures. The fact that the rising number of repetitive motion injuries is associated with individuals working with computers indicates that this analysis is not just a handy metaphor. Shifting the control to the eye did not bring any benefits either. Besides its physiological role in finding and interpreting visual clues, in the traditional GUI the eye is forced to play the role of the white cane of the blind for the hand. All of the cursor guidance and positioning, often demanding a single pixel precision (in graphic programs), is done under the guidance of the eye. This leads to overstraining of this sensory channel to the point that computer operators 'forget' to blink, ultimately developing the chronic dry conjunctivitis of the eyes.

In our other daily activities we rarely depend on single-sense feedback. Take, for example, the simple act of putting a pencil on the table. Although the eyes are involved, their role is more general – seeing whether the surface of the table is within our reach, and if it is clear of other objects. The actual act of putting the pen on the surface is guided more by cutaneous and proprioceptive clues, and augmented by discrete but definite auditory cues. This multimodal and complementary feedback is the reason our daily actions do not cause over-straining of any particular sense.

Introducing Multimodal Interaction

Although the term multimodal interaction encompasses both input and output, I will focus more on ways of making the output in human-computer interaction more tangible by using different sensory channels. This should not be taken as an indication that multimodal input (for example, using both speech and gestures) is of lesser importance. It just introduces another level of complexity that goes beyond the topic of this paper.

The value of multimodal output (feedback) was recognized in HCI design and is currently most widely-

187

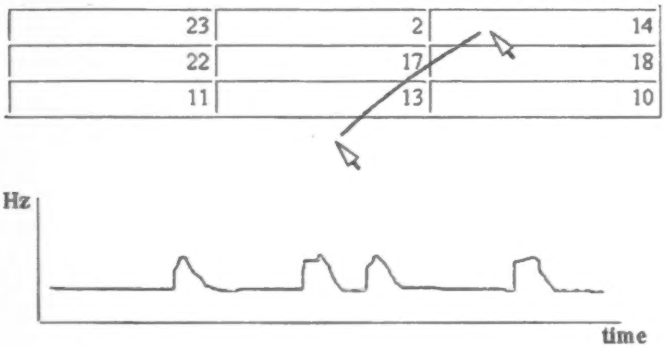

Figure 4. Movement of the iFeel mouse (cursor) over the table cells is augmented by haptic output produced by vibro-tactile unit in the mouse. Although the stimulus is fairly discreet it significantly reduces the load of the visual sense during table input or selections using pull-down menus.

spread as coupling of actions to sounds. Sound production is a standard part of modern computer systems, and it is astonishing that the value of consistent sound feedback was not recognized earlier and integrated into interface design guidelines. By coupling actions to sounds I do not mean often exotic "sound schemes" featuring drum-rolls for window closings and alarm clock sounds for warning messages. An example of consistent auditory feedback is a discrete "click" associated with opening of a new window in Microsoft Internet Explorer. The sound is so discrete that many users don't consciously perceive it, yet immediately notice its absence when they switch to another browser.

Another indicator of this trend is the introduction of a mouse (Logitech, iFeel ™ mouse) that allows the user to "feel" different objects (for example, folders) and actions (dragging, scrolling) in the traditional GUI interactions. Haptic feedback is provided by a vibro-tactile unit in the mouse and can be finely tuned to fit individual preferences.Although the additional information initially seems trivial and meager (a series of vibration patterns), it becomes very quickly evident that it significantly increases the comfort of interactions (Figure 4.).With complementary information about cursor location or action, the user does not have to rely as much on an already over strained sense of vision.The difference in the experience can be compared to the difference between typing on a standard keyboard where every key-press is accompanied by a tactile, kinesthetic and auditory cue, and typing using keyboards where the keys are outlined on a touch-sensitive surface and provide no specific feedback. Numbers seem to confirm the benefits to the user of haptic technologies – in the first year they were introduced Logitech sold a quarter of a million of iFeel™ mice (quoted from Immersion TouchSense™ Web site).

Manifest increase in the quality of interaction experience is making the devices that use haptic or force-feedback a standard in the area of computer games.A quote from Briggs and Srinivasan illustrates the use of haptic interfaces in PC game playing:

> Active haptic interfaces can improve a user's sense of presence: Haptic interfaces with 2 or fewer actuated degrees of freedom are now mass-produced for playing PC videogames, making them relatively cheap (about US$100 at the time of this writing), reliable, and easy to program. Although the complexity of the cues they can display is limited, they are surprisingly effective communicators. For example, if the joystick is vibrated when a player crosses a bridge (to simulate driving over planks) it can provide a landmark for navigation, and signal the vehicle's speed (vibration frequency) and weight (vibration amplitude). (Briggs and Srinivasan, 2001)

An important area where multimodal sensory feedback plays a crucial role is the area of affective com-

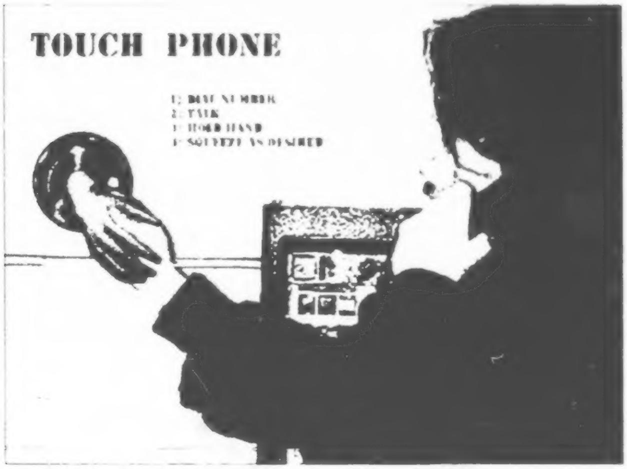

Figure 5. By "holding hands" one can carry out both verbal and non-verbal parts of a conversation. It remains to be seen whether the actual emotional "value" to the users is high enough for this invention to be economically viable. (photo and hand simulation by S. Milekic, modeled by M. Lengauer)

puting. The pioneer of research in affective computing, Rosalind Picard (Picard, 1997) suggests a number of possible applications where the affective state of a human user becomes accessible to a computer, or another remote user. One of the applications is TouchPhone, developed by Jocelyn Sheirer (described in Picard, 2000), where the pressure that a participant in a phone conversation applies to the headset is transmitted to other party's computer screen as a color range - blue corresponding to slight pressure and red corresponding to the maximum pressure value. Inspired by TouchPhone I dared imagine a more sophisticated model where haptic information is transmitted back and forth by holding the other party's simulated "hand" while carrying on the conversation (Figure 5.).

Besides making abstract data manipulations tangible and accessible to humans and providing a channel for affective communication, using complementary sensory feedback to illustrate complex physical interactions is becoming a method of choice, especially for getting feedback while operating complex machinery or vehicles. The paradox here is that the human operator is most often not directly exposed to the relevant physical changes, and these are made available by *translating* numerical data into a sensory experience readily interpretable by a human. The value of adding additional *simulated* sensory information to the real world task is beautifully illustrated by an example provided by cognitive scientist David Kirsh:

> This odd situation which digital technology creates is nicely portrayed by the way modern airplanes rely on simulations of the feel of flying to improve the control of pilots. Apparently, jets fly faster if their center of mass is moved closer to the plane's nose, thereby changing the relative position of the center of mass with respect to the center of lift. The trouble is that in moving the center of mass forward there is an increased chance that the plane will tip into a nose dive. To keep the plane flying on this knife edge the speed and sensitivity of adjustments is so great that pilots can no longer use mechanical means to control their planes. To assure fast enough response such jets now rely on digital networks to relay a simulated feel to the pilots. When a pilot pulls up on his steering wheel the computers inside the plane simulate the resistance of the ailerons delivering to the pilot the haptic information he or she needs to know what they are doing. Small computer adjustments augment and speed up these pilot reactions. To the pilot this force feedback is an integral

BEST COPY AVAILABLE

> part of the way he or she flies the plane. But, of course, there is no true resistance in the steering wheel. Pulling harder on the wheel is just a way of sending the number 7 to the wing actuators instead of the number 5. Computation is so irremediably built into planes that pilots could be in simulators. (Kirsh, 2001)

There is yet another area, academically not that well researched, that readily embraced (no pun intended) the prospect of multisensory interaction. This is the vast domain of on-line sex. As NBC reports, there exist already a number of (multisensory) products that belong to the new "cyberdildonics" area, as well as full "cyber sex suits" that allow a wide variety of tactile sensations to be experienced on strategic body parts (Brunker, MSN NBC online).

Barriers

While it may seem that all arguments are in favor of building tangialities, it is worth investigating the barriers and problems associated with this approach. As I see them, there are some closely related general problems that are listed here separately only for the sake of clarity. These are:

- success of "drag-and-drop" and "point-and-click" interface;
- failure to realize that changing to, or addition of, another interaction device calls for redesign of the GUI;
- need to un-learn established (imposed) conventions.

Although it may seem to be a paradox that an early solution would be blocking the introduction of more advanced ones, historically this is a common occurrence. The QWERTY keyboard arrangement became standard for typewriters, and continued to be used even for computer keyboards where any conceivable keyboard layout is equally accessible. In the same way, the very success of mouse/cursor point-and-click interface made it a de facto standard closely associated with the very concept of what it means to interact with a computer. This standardized notion of what an interface should look like affected adversely the introduction of any new interaction device. For example, using a touchscreen as an interaction device while preserving the traditional GUI creates huge problems for the user because of the discrepancy in scale of objects necessary for comfortable interaction. An icon with dimension of 8 x 8 pixels; for example, a window closing box, is perfectly acceptable for a single pixel active tip of the cursor, but is definitely inappropriate for finger interaction. In environments adequately scaled for finger interactions, touchscreens have been shown to be superior pointing devices (Sears, Shneiderman, 1991). The same holds true for other interaction devices - introduction of continuous speech recognition calls for adequate feedback that verbal information has been successfully transmitted. The use of haptic mice and joysticks introduces the texture as a GUI design element, etc.

Another general problem that has to be taken into account when introducing new ways of interaction is the need for un-learning of adopted conventions. This can be a very slow process, and a transitional stage should be a part of the design of any new convention (and lack thereof is often the reason for their failure).

More specific problems with the introduction of multimodal feedback come from our lack of knowledge of the complexities of multimodal interaction. Just adding another channel to human-computer interaction is not by default beneficial. A logical and commonsensical analysis tells us that additional information presented through another channel (like any other information) may fall into the following categories:

- conflicting
- competing
- redundant
- complementary

In the above list, only the last category has a beneficial effect on interaction. In the following paragraphs I will provide examples of different types of multimodal information.

Conflicting

Conflicting information is often the consequence of hardware/bandwidth limitations as is the case with lack of audio/video synchronization in streamed Web- videos. Sometimes it is a product of poor design; for example, when an animated character's mouth movements are inappropriate for actual utterances. Conflicting information is sometimes purposefully designed into an application. An example is some Web site designs which try to keep the user "glued" to the site in order to artificially boost ratings.

Competing

An example of *competing* multimodal information is a voice overlay that is not synchronized to the printed text one is trying to read. Animated graphics (Gifs) on Web sites are another example of information competing for the same sensory modality (visual) and claiming a part of cognitive resources. An example from everyday life is the effect carrying on a phone conversation has on driving ability.

Redundant

A definition of *redundant* multimodal information conveyed through another sensory channel, which does not increase the total amount of information about the interaction but is also not adversely affecting the interaction.

Complementary

Complementary multimodal information is information conveyed through another sensory channel that does increase the total amount of information received and has a beneficial effect on interaction. This effect can be manifest as an increase in efficacy of interaction, or decrease in number of errors. This is the only instance where the bandwidth of human-computer information channel is increased by engagement of another channel.

Case for Building Tangialities

In conclusion, one can make the case for building tangialities for the following reasons:

- widening bandwidth of human/computer communication channel;
- adding affective dimension to interaction;
- allowing grasping and manipulation of complex concepts without the need for explicit formalization;
- reducing cognitive load by use of intuitive body (biological) knowledge;
- reducing the strain on one sense (vision) - single sense fatigue;
- possibility of adding another dimension to meta data – "how does it feel" (being able to feel the texture of paintings and other, otherwise "untouchable" objects)

Another significant use of tangible descriptions of data and results of data manipulations is in making the digital domain more accessible for populations with special needs. There are already some promising results in this area (Yu, Ramloll and Brewster, 2001; Gouzman, Karasin, Braunstein 2000).

Currently, we are lacking a satisfactory theory of multimodal interaction and most often arrive at usable results through a process of trial and error. It is evident that this theory has to come from interdisciplinary efforts bridging the disciplines as diverse as neurophysiology, tele-robotics and computer science.

References

Biggs, J., Srinivasan, M.A. (2001) Haptic Interfaces, in Stanney, K.M. (ed.) *Handbook of Virtual Environment Technology*, Lawrence Erlbaum Associates, Inc.

Brewster, S.A.(2001). The Impact of Haptic 'Touching' Technology on Cultural Applications In Proceedings of EVA 2001. (Glasgow, UK), Vasari UK, s28 pp1-14, also available on line: http://www.dcs.gla.ac.uk/~stephen/papers/EVA2001.pdf

Brunker, M., (2001). Sex toys blaze tactile trail on Net, MSNBC News online article:http://www.msnbc.com/news/318124.asp?cp1=1

Gouzman, R., Karasin, I., Braunstein, A. (2000) The Virtual Touch System by VirTouch Ltd: Opening New Computer Windows Graphically for the Blind, Proceedings of "Technology and Persons with Disabilities" Conference, Los Angeles March 20-25, 2000, on-line: http://www.csun.edu/cod/conf2000/proceedings/0177Gouzman.html

Immersion TouchSense™ technology, consulted on line 24/3/02: http://www.immersion.com/products/ce/generalmice.shtml

Kirsh, D., (2001) *Changing the Rules: Architecture in the new Millennium*, Convergence

Logitech, http://www.logitech.com

McCullough, M. (1996) *Abstracting Craft: The Practiced Digital Hand*, The MIT Press, Cambridge, MA

Merriam-Webster Dictionary consulted on line 2/28/02: http://www.m-w.com/home.htm

Norman, D. (1991) Cognitive Artifacts, in Carroll, M.J.(ed.) *Designing Interaction: Psychology at the Human-Computer Interface*, Cambridge University Press

Oakley, I., Brewster S.A. and Gray, P.D. (2000). Communicating with feeling. In Proceedings of the First Workshop on Haptic Human-Computer Interaction, 17-21

Patten, J., Ishii, H., Hines, J., Pangaro, G. (2001) Sensetable: a wireless object tracking platform for tangible user interfaces, in Proceedings of the SIGCHI conference on Human factors in computing systems 2001, Seattle, Washington 253-260, ACM Press, NY

Picard, R.W. (1997) *Affective Computing*, The MIT Press, Cambridge, MA

Picard, R.W. (2000) Toward computers that recognize and respond to user emotion, *IBM Systems Journal*, Vol. 39, 705-717

Roy, D. (2000) Learning from multimodal observations. Proc. IEEE Int. Conf. Multimedia and Expo (ICME), New York, NY, consulted on line 3/4/02: http://dkroy.www.media.mit.edu/people/dkroy/papers/pdf/ieee_multimedia2000.pdf

Sears, A., Shneiderman, B., (1991) High precision touchscreens: design strategies and comparisons with a mouse, *International Journal of Man-Machine Studies* 34, 593-613

Shneiderman, B. (1983), Direct Manipulation: A step beyond programming languages, *IEEE Computer*, 16, 8, 57-69

Shneiderman, B. (1998) *Designing the User Interface: Strategies for Effective Human-Computer Interaction*, Addison-Wesley

Ullmer, B., Ishii, H. (2000) Emerging Frameworks for Tangible User Interfaces, *IBM Systems Journal*, v393n3, 915-931

Yu, W., Ramloll, R., Brewster S.A. (2001) Haptic graphs for blind computer users, in Brewster, S.A. and Murray-Smith, R. (Eds.) *Haptic Human-Computer Interaction*, Springer LNCS, Vol 2058, 41-51.

Making It Realtime: Exploring the Use of Optimized Realtime Environments for Historical Simulation and Education

Chris Calef, Turlif Vilbrandt, Mythworks, USA/Japan, Carl Vilbrandt, University of Aizu, Japan, Janet Goodwin, Aizu History Project, USA/Japan, James Goodwin, University of California, USA

http://www.myth-works.com/simtools and http://www.u-aizu.ac.jp/~vilb/aizu_history

Abstract

As museums and educators struggle with the challenges of presenting their material in a digital format, many overlook the application that has spearheaded the development of virtual reality for the average consumer: 3D realtime game engines. These 3D game engines offer greater versatility, usability, maturity, simulation and codebase than most current 3D realtime frameworks. At the University of Aizu, we are using the Quake engine in conjunction with the Povray raytracing engine to attack the problem of visualization and simulation from two sides. We have modeled a temple from northern Japan that users can experience in realtime. However, to deal with the limitations of simulation in realtime, we have added the ability for users to select a view for greater detail. The selected view is rendered in the background as the users continue to travel through the temple, and is delivered in a separate window when finished. Our paper will describe relevant game paradigms, their usefulness, and our work in detail, including problems and solutions we have discovered along the way, and conclude with suggestions on how this work could assist museums and educators in simulation and modeling.

Keywords: Simulated Environments, Virtual Reality, Realtime, 3D Games, Cultural Heritage, Historical Restoration, Virtual Museum, Quake, Id Software, Povray

Introduction

The Web offers unprecedented opportunities for museums to escape the physical confines of their buildings and reach a vast new audience. Many institutions have begun to take advantage of this by using VR and 3D technologies, but this use does not take full advantage of the virtual information world we are entering. One factor to be aware of is the power of realtime modeling and simulation. This is useful both for creating engaging on-site exhibits, like the popular "virtual fish tanks" (MIT, 1999), and for reaching across the internet with immersive, educational simulation programs. There are widely varied approaches to realtime 3D modeling and simulation, but this paper will focus on game engines and solutions. Typical 3D environments used for academic and historical purposes are often hard to navigate, obtuse and short on interactivity. But 3D games, on the other hand, have been designed from the ground up to be usable, enjoyable and very interactive. The 3D gaming environments typically go further than just modeling a static object or environment, but instead try to simulate some additional properties and interactions. People in the academic community should embrace and extend gaming concepts and technology as a means of simulating, storing, testing, and transmitting their ideas.

3D Gaming = Usability = Learning

Taking advantage of realtime 3D game paradigms yields several advantages: increased user enjoyment, increased use of the application, and transparent learning. Games have been designed from the ground up for usability and fun. The more hours a user spends in a game environment, the better it tends to do in the market place, and the more money it makes. As a result, the primary focus for most game companies is on making 3D environments that are highly functional, easy to learn, and enjoyable to use.

By embracing game code and techniques, a planned application can instantly come up to speed with a usable, sophisticated interface, in an environment that has proven staying power. Users come back to their favorite games again and again. Game techniques in conjunction with modeling and simulation can yield a very interesting potential byproduct for

BEST COPY AVAILABLE

academic applications: transparent learning. If users are constantly interacting with a program for enjoyment, they will pick up a variety of skills and knowledge without approaching it as a learning experience, and in some cases learn without even realizing it. Currently, Mythworks and the Oregon Center for Applied Science are working under a grant for the National Institute of Health on a project that involves teaching children how to navigate and cross streets safely. One of the goals of the program is to make children feel as though they are playing a game, allowing the skills to be learned through modeling and simulation. The more a child, or any user, comes back to such an environment, the more the modeling is reinforced and the more the skills become "second nature". Making sure an environment is "playable" goes a long way toward this goal. This has also been demonstrated in flight-simulator-based games in which users who fly fighter jets in air combat and other such engaging simulations demonstrate a highly accelerated learning curve when learning to fly a real plane (Hampton, 1994). There is at times a balance that must be struck for an application between accurate modeling and playability, but there is no doubt that the tools and techniques developed by the game industry hold great promise for academics and educators.

A Tested Interface for Free

One of the most important elements to be gleaned from the gaming community is the set of user interfaces that have evolved for control of player movement and view direction. In contrast to other virtual environments, the typical 3D game has evolved as an arena for competition between players. This means there is no room for anything but the most efficient interface between the player and the computer. In VRML, on the other hand, the single worst feature is its viewer interface. Most VRML players are difficult at best to navigate in; usually only the mouse is used, often with a very counterintuitive set of controls. In Cosmo, for instance, one must grab the center of the screen and then stretch a line out from it to move in a certain direction. Compare this to the very natural and efficient command standard used in Quake (ID, 1996) gaming, in which the mouse determines "look" direction and buttons or keys apply forward, backward, and sideways motion. The rest of the interface choices in Cosmo and other VRML clients just get worse, and there is usually no easy way to modify them. Also, Quake and other game engines offer support for additional human interfaces like joysticks. Although in some VRML players it is possible to program a new interface, this can be time-consuming and difficult.

Perhaps the most important lesson to be learned from gaming interfaces is that many of these choices have been made not by the gaming companies themselves but by the users, over long periods of trial and error. Game companies learned long ago to leave interface choices up to the user, and as a result the users have found the best combinations for different types of games, goals, hardware interfaces, and handicaps. This natural evolution of 3D navigation should not be ignored, and as the gaming companies learned, it should always be easy to change.

Making It a Sim

Yet another aspect of many 3D games is representation and simulation of the natural world. Nearly all such games implement basic Newtonian forces like gravity. Other games go much further. In Black & White, by Lionhead Studios Ltd., the user plays the role of a god who rules over the population and resources of a small island. The game simulates population growth and decline, natural disasters, disease, and social interaction. Weather in the game has an impact on the growth of vegetation and crops. Going even further, if a user is connected to the Internet, the game can actually check the player's online local weather report and simulate these conditions within the game environment.

While most museums have yet to discover 3D gaming and simulation techniques, in the past few years many have discovered the utility of "raytracing" programs for visualization. Applications like 3D Studio Max, Maya, Poser (for character animation), and many others have provided the capability for modeling virtual environments, enabling 3D "walk through" animations of historical sites, graphic representation of scanned artifacts, and so forth. Although there are too many examples to begin to list them all, one of the most impressive of these "walk throughs" was made by the "Virtual Olympia" project (Ogleby, 1999) directed by Cliff Ogleby of the University of Melbourne.

While raytracing software can be very useful for 3D visualization, we find that on their own, these

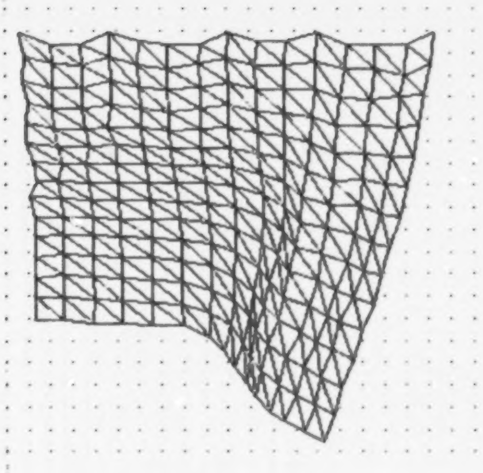

Fig. 1: Surface Representation

applications lack many features that would be necessary for us to use them for modeling and simulation. Most commercial and open-source raytracing environments, and 3D games for that matter, focus exclusively on representation of surfaces. While surface representation can be adequate for creating relatively static scenes, more tools and better data sets are necessary in order to accurately portray a dynamic world. A good example is an animator who wishes to model a scene involving the eruption of a mountain, with rocks flying into the air. Using old-style raytracing software, the animator would have been required to "fake" the paths of the rocks, by defining arbitrary curves for them to follow. This was because the rocks were surface representations only, and as such even if the environment included gravity the rocks would have no mass for it to affect. Most packages nowadays do better than this, implementing some procedural tricks with the surfaces to give the impression of a gravity algorithm. However, these tricks are extending surface data and math beyond its natural and practical limits. Because the code is based on mathematical "tricks" it can be implemented in many different ways. This allows the primary commercial packages to each come up with their own method of handling and storing object or material properties, such as density, tensile strength, and mass distribution. This makes it impossible to share this kind of data between applications, or even within a single software package. Without this "real" mathematical data, accurate simulation is impossible, and without a shared format, the user will find it impossible to adequately combine models from more than one application. Most 3D games share these portability and accuracy issues.

As an example, consider an animator who is using Poser to do cloth simulation on a walking human figure, and then needs that human figure to brush against a tree created in PlantStudio. She will likely find that while the cloth has a collision detection algorithm it uses to respond accurately to the human figure, there is no such algorithm to enable it to interact with a model from another program. Furthermore, the cloth algorithm probably does not include the possibility of the fabric catching on a branch and tearing, because this situation does not happen in a software package focused only on human figures. In order to have an interaction between models, there is a need for an open object library that includes simulation logic, so that cloth, humans, and trees can be handled within the same procedural framework.

SEDRIS (Sedris, 2001) seems to be the best hope for an implementation of a common format for simulation, and we are excited to see how it develops. Many vendors readily support OpenFlight and SEDRIS formats. We have questions regarding the depth and breadth of this support, however. To the best of our knowledge, there has been no attempt to apply SEDRIS logic to any game framework except OpenFlight, and as part of our research we hope to make some contribution toward linking SEDRIS to the Quake engine.

Simulation logic can help an animator with tasks far beyond simple physical simulations like cloth behavior and falling rocks. Artificial intelligence for moving "actors" in the scene is one prime example. The value of introducing AI into a raytracing environment was proven dramatically in the recent film "Lord of the Rings: Fellowship of the Ring" with the use of the MASSIVE simulator program, which automated much of the combat AI for the movie. This enabled the director to create gigantic battle scenes that would have otherwise required a prohibitive number of animator hours. The same concept could be applied to more peaceful purposes, enabling animators to automate the life of an entire village, for instance, based on a few simple behavioral rules for each actor.

For this kind of simulation to be really effective, a cross-disciplinary approach is necessary. For ex-

Fig. 2: AI "Bots" in Quake

ample, with a historical simulation, only part of the "rule set" will fall within the bounds of what we would ordinarily call "history", and the rest will fall into other academic fields, such as physics, architecture, geography, and botany. In order to provide a useful three-dimensional simulation of a Japanese village, for example, there would first need to be historical knowledge about what kind of people lived there, what crops they raised, what kind of buildings they lived in, and so forth. To make the simulation work, there would also need to be a set of rules describing how fast an object falls if someone drops it, what happens when it hits the ground, and so forth — clearly the realm of the physicist. To animate the people of the village is also outside the traditional domain of the history department and belongs, rather, to the field of biomechanics. Meanwhile, the land area around the village should not be empty space; it should reflect the topography and ecosystems one would expect at this particular time and place. This calls for expertise from geographers and botanists. Meteorologists could contribute a working weather model, which would have a direct impact on the agricultural cycle.

The possibilities for simulation are as infinite as the complexity of the natural world. The choice for educators, then, becomes one of deciding, for a particular application, what form the simulation should take and what aspects of reality should be included.

Realtime versus Rendertime

In addition to their enormous utility for the animator wishing to create realistic still scenes and animations, simulation tools also make possible an entirely different kind of application. In realtime simulation using gaming methods, a user can interact with the program, as well as with other users, in an immersive and entertaining environment. This ability does not come without cost, however. In offline simulation, the only limitations to detail are the accuracy of the algorithms employed, and to a lesser extent the computing time necessary to reach the desired degree of accuracy. If necessary, the simulation can be allowed to run for days or weeks to attain this accuracy. In contrast, the restrictions imposed on a realtime environment are significant. Even at a relatively slow frame rate of ten frames per second (barely adequate for games), all computation for each frame must be finished within 100 milliseconds. Even on the most powerful gaming systems, this hardly allows for unlimited complexity in the simulation model. Many CPU-intensive algorithms are simply impossible to model in this environment. In many cases, processes that are too difficult and time-consuming for true realtime computation can be precalculated and rendered into prerecorded animations or stills, which are then sent to the user as a movie or used inside the realtime game environment, transparent to the user. Using a server-client model we can have real-time interactivity and simulation on a low power client machine, backed up by super computers or distributed computing to produce high detail simulations, images and animations.

Currently at the University of Aizu, we have modeled such a system to run on a single computer using Quake as the front end "client" with a thin bridge to Povray as the backend "server". Our test case is a model of Enichiji temple, from the Aizu region of northern Japan. (See (Vilbrandt, 2001) for an early version of the Enichiji model.) In order to provide an immersive environment, we have created a model of this temple which runs in Quake, and allows the "player" to climb the stairs, inspect the internal architecture, and move under, over, through, or around the temple in full realtime.

However, to overcome the unavoidable limitations on complexity of a realtime application, the "player" may at any point choose a scene to be viewed in greater detail. Using Povray, a much more complex model of the chosen scene is rendered, creating still images or animation to be viewed in a separate window. We are working toward extending the complexity of this scene, so that in addition to the temple, there will be a section of terrain, flags blow-

Fig. 3: Inside Temple, with Flag

ing in the wind, various vegetation and rocks, an animated monk figure, and a flock of crows. This will give us a better test case for demonstration, because the interactions among the wind, the flags, the birds, and the monk's robes will be impossible to render in Quake at the same level of detail that would be possible in Povray.

Another way we have discovered to combine "rendertime" raytracing with "realtime" gaming is by precalculating special effects to be included in the realtime game. As a trivial example, imagine stirring cream into a cup of black coffee. The cloud shapes that swirl around the cup could be simulated using a particle-based fluid motion algorithm, but this is almost certainly too much computation to run in a realtime gaming environment for such a minor effect unless this action is for some reason critical to the game. Instead, the motion could be simulated in advance and stored as one of the animations possible for a "cup of coffee" object, and unless the user tried to pour the cream from the other side of the cup, he or she would likely never know the difference.

In our research, we found this technique of precalculating simulations to be mandatory for presenting any kind of adequate realtime simulation. Cloth animation is a good example; in our test model, a hanging cloth was draped in a doorway of the temple, and a wind effect simulation caused it to blow back and then settle down again. The cloth was first simulated using a particle algorithm in C++. It was then written out as a series of frames in Povray include files. When a suitable sequence of frames was found, after adjusting the parameters of the simulation to get the most realistic movement, then the coordinates of the cloth mesh vertices were imported to a Quake model format. After that, the same animation could be called from within Quake whenever game logic demanded.

Current Tools, Work and Web Interface

For an offline rendering application, we chose to use Povray (Hallam, 1995), an open source raytracing program with a number of features that we found useful. It has a fairly usable scripting language, but of even greater value was the ability to call an external application between frames. We use this option to call our executable program, written in C++, with the frame number and animation clock sent as arguments. The C++ code then handles all moving entities, physics, collision detection, etc., and after determining the new position and orientation for each visible entity in the current frame, writes out a POV script file, which is then rendered.

For our realtime game engine, we are using a modified version of the Quake game engine, released by Id Software under the GPL license. While a game engine has drawbacks in terms of supporting limited platforms and requiring users to download a piece of software, we feel that these limitations are more than balanced out by the speed, realism, overall versatility and extendibility offered by such a solution. Under GPL, any historical or educational game created with the Quake engine can be given away for free or sold for a profit, providing only that the source code is made available to the public. For academia, this should be a plus.

We have specifically chosen the open-source Quake game engine because it is one of the most-used 3D game engines. The engine is fast and small. It was designed 5 years ago for pentium class machines and therefore has a broad base of systems on which it can run. However, due to development by the open user community, this version of Quake has evolved to take advantage of new hardware and software techniques and has begun to rival even the latest Quake 3 engine from Id Software. If a project based on the Quake engine is managed and engineered properly, it can have both low-end compatibility and high-end glitz and hardware optimization in the same application. In addition, Quake runs on ALL the major desktop operating systems, in-

Fig. 4: Landscape with Trees

cluding Linux. Most importantly, both Quake and Povray have Free Software licenses allowing easy modification by anyone, and solving some serious problems associated with data transparency. This means three things: one, that we have a proven 3D code base to work with; two, that our work and that of others after us can be preserved; and three, that persons or institutions will be able to make adjustments or modifications of our work in the future without need of our presence or permission.

In the process of our research, we have also developed a library of C++ code which is available from our Web site. Some of this code would be redundant to anyone already using a high-level simulation system such as SEDRIS, but some of it may have relevance to people wishing to experiment with procedural approaches to Povray and/or Quake. Among the possibly interesting functions are:

- Landscape generation — using the simplest possible application of the diamond-square algorithm.
- Buildings — procedural generation of some very simple building shapes. Not accurate enough for a detailed model, but perhaps useful for low-LOD background structures.
- Human skeletal animation system — takes input from motion capture, .data files, and quake .mdl files.
- Basic physics code — including gravity, drag, acceleration, collision detection.
- Cloth simulation — with limited collision detection against a mesh.
- Webinterface tools — Some of our code uses Postgres as a backend database server for storing and sharing simulation data. We have a developed an advanced html interface to manipulate and manage this database over the Web.

All of our code writes out to both Povray and Quake, in some form or another, unless it would be redundant in one of the applications. In Povray, most objects are written described as triangle meshes, whereas in Quake the format depends on the type of entity in question. Mobile, animated entities (humans, animals, etc.) are written as .mdl files, which contain a description of polygons, a set of animation frames, and a "skin" texture for the model. Larger, static entities such as buildings are written as BSP "brush" entities (see Feldman, 1997). Of course, for the same scene to function in both applications, the coordinate system must be identical. For our project, that meant adopting Quake coordinates to Povray.

Problems Unsolved

We have found that one of the major limitations in using Quake has been BSP (Binary Space Partition) representation of environments. The use of BSP trees enables the game engine to sort the entire scene into "potentially visible sets" of polygons. (For a more complete description of the BSP algorithm, see Feldman, 1997.) While this works just fine for the simple concave environments used in Quake (where the environment consists entirely of rooms linked by tunnels), it does not work so well for complicated and convex objects and environments. The worst-case example so far has been the compound convex shapes found in the roofs of Asian architecture. Some convex roof examples are acceptable while others produce such complicated BSP trees that they become unusable. Another issue with BSP trees is they are not easily modifiable in realtime, making them difficult or impossible for some aspects of simulation.

Another difficulty with the Quake engine is its inability to support large open terrain areas, making scene design very challenging when one is attempting to simulate an outdoor environment. It is usually necessary to artificially "box in" a scene using terrain entities such as a "wall" of forest, a hedge, a rock wall, or some other device to limit the potentially visible area from any given point in the scene. When too much complexity is visible at any given

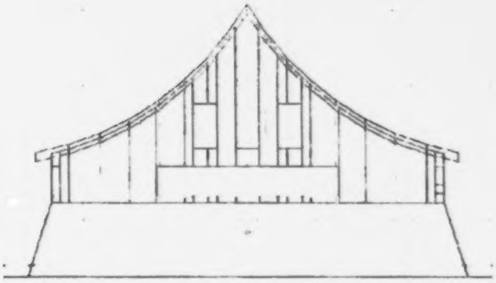

Fig. 5: Wireframe of Autogenerated Building

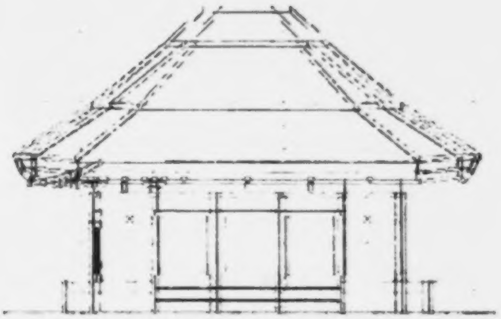

Fig. 6: Example of a Difficult Roof

time, the game can error by "graying out" whole sections of the scene, interfering dramatically with the overall sense of realism.

Finally, the game can handle only a limited number of moving entities in a scene, like human figures, animals, etc. The exact number is dependent on the hardware configuration and the amount of AI being run for each entity, but the limit seems to be something on the order of 40 to 60 entities visible at one time. This is enough for most simulation purposes, but it can pose significant limitations on scenes involving large crowds of people, herds of animals, and so forth.

Future Work

This work is in its infancy, and as such currently has very little simulation support outside of basic physics and collision detection. In addition to adding links to SEDRIS, we have plans to incorporate a much improved skeletal animation system utilizing a genetic algorithm to minimize energy expenditure. (See the work of Schmitt and Kondoh, in Schmitt, 2000). We also intend to increase our use of particle and voxel systems for solids representation, and to add links to the HyperFun library (Pascoe, 1996) for function representation of solids (F-Rep).

Our main focus in the near future will be to extend our current Quake-to-Povray framework. We intend to fully implement a realtime Quake client communicating with a backend Povray server across the Web. Users will be able to choose a scene or path through the realtime environment and have it rendered and returned to them in high detail. This will allow institutions to offer the equivalent of mainframe processing power to the average home computer user.

Conclusion

We see the combination of simulation tools with realtime, potentially multiplayer gaming to provide museums and academics with any number of new ways to involve visitors in interactive learning experiences. Also it has been demonstrated that an effective application can create a community around it. This means larger and lasting participation in given fields, exhibitions or focus and for museums it means more visitors. Some applications might include the following:

- Users could participate in an industrial process, running an airplane or automobile factory, or being part of the operation of an early coal-fired electric plant.
- Museums could have "game rooms" with many computers networked together, allowing visitors to take part in multiplayer interactive simulations.
- A school class could become the population of a farming village, and spend the afternoon planting wheat, learning to fix sheds and houses using appropriate tools and resources, deciding what crops to plant, where to clear forests, where to trade and what to barter for. They could learn firsthand the need for pottery, because when the villagers stack the grain in open piles or in sheds, the water comes in and their next year's supply of food rots. Accurate simulation could require the players to build a kiln hot enough to fire the clay that the villagers dig up nearby, and that would help determine the amount of wood that the village harvests. This sort of game can teach constantly without the participants ever even becoming aware of the instruction.

- In an astronomy exhibit, visitors could view the orbits of the planets around the sun, or stars around the galactic core. Users could navigate a virtual spaceship or modify the masses of the stars and planets and observe the forces of gravity.

The content in these facilities could be updated regularly, for very little cost, by museum staff using an html interface to the simulation database. Constantly changing content could keep people interested and returning to the site or to the institution to see what is currently "going on" in the simulation. Any of these projects could also be made available over the general Internet. For museums, this means allowing people anywhere in the world to have access to the information stewarded by the museum, and as a byproduct potentially increasing membership.

As a final note, much of the logic and work done in the field of modeling and simulation seems to be related either to violent video games or military applications. We wish to be part of a move toward the exploration of more peaceful and educational subjects for simulation. Considering the dangers currently faced by heritage sites and natural resources as a result of human war, overpopulation, and over-consumption, it can only be a good thing for museums and public educators to have access to better tools for reaching out to the general public. Using realtime game engines can create immersive and entertaining environments for educating members of the public about the processes and forces that impact our lives, our history, and our future.

Acknowledgements

This work would not have been possible without the support of Drs. Naotoshi and Yuko Takeda of Shiokawa, Fukushima-ken, Japan; Mr. Kyoichi Okubo of the Aizu Digital Valley Promotion Association of Shiokawa, Fukushima-ken, Japan; Ms. Miyako Aotsu of Aizu-Wakamatsu, Fukushima-ken, Japan; Id Software and the online Quake community; the Povray team and online user group; and the Free / Open Source Software movement.

References

Feldman, M. (1997). Introduction to Binary Space Partition Trees. Last updated 25-Aug-1997. http://www.geocities.com/siliconvalley/2151/bsp.html

Free Software Foundation. (1985). Last updated 30-July-2001. http://www.gnu.org/fsf/fsf.html

Hallam Oaks Pty. Ltd. (1995). POVRay. Last accessed 20-Feb-2002. http://www.povray.org

Hampton, S., Moroney, W., Kirton, T. & Biers, D. (1994). The use of personal computer-based training devices in teaching instrument flying: A comparative study. Daytona Beach, US: Embry Riddle Aeronautical University

Id Software. (1996). Last updated 29-Oct-2001. http://www.idsoftware.com

MIT Media Lab & Nearlife. (1999). The Virtual Fish Tank. Last updated 03-Dec-200. http://www.mos.org/exhibits/current_exhibits/virtualfishtank

Ogleby, C. (1999). Virtual Olympia. Last updated 03-Nov-2001. http://www.phm.gov.au/ancient_greek_olympics

Pasko, A. (1996). HyperFun Project. Last updated 19 -Dec-2001. http://www.hyperfun.org

Schmitt, L.M. & Kondoh, T. (2000). Optimization of Mass Distribution in Articulated Figures with Genetic Algorithms. Aizu, Japan: University of Aizu

SEDRIS. (2001). Synthetic Environment Data Representation and Interchange Specification. Last updated 03-May-1998. http://www.sedris.org

Vilbrandt, C., Goodwin, J.M. & Goodwin, J.R. (2001). Digital Digging: Computer Models of Archeological Sites:: Enichiji in Aizu, Japan. *2001 PNC Pacific Neighborhood Consortium Annual Conference and Joint Meetings*. Taiwan: Computing Centre, Academia Sinica

Networked Multi-sensory Experiences: Beyond Browsers on the Web and in the Museum

**Fabian Wagmister and Jeff Burke, HyperMedia Studio,
School of Theater, Film and Television, University of California,
Los Angeles, USA**

Abstract

The defining characteristic of the digital era is the potential that it brings for "real-time" interconnection between anything that can be measured, expressed, or controlled digitally. The World Wide Web stems from one type of digital interconnection: well-defined standards linking a "browser" with remote machines presenting information to be browsed. Yet digital technology enables more than just new approaches to presentation, browsing, and searching. It can create dynamic connections between different physical spaces and across sensory boundaries, and provide experiential interfaces for interaction that move beyond the mouse, keyboard, and screen. It can relate the physical space of the museum to the virtual space of the Web for both individual and group experiences.

Using our past media-rich installation and performance work as a reference point, this paper will present a vision of digital technology for the museum as a dynamic connection-making tool that defines new genres and enables new experiences of existing works. The authors' recent works include the interactive media-rich installations *Time&Time Again…* (with Lynn Hershman, commissioned for the Wilhelm Lehmbruck Museum, Germany), *Invocation & Interference* (premiered at the Festival International d'Arts Multimédia Urbains, France), *Behind the Bars* (premiered at the Central American Film and Video Festival, Nicaragua), and *…two, three, many Guevaras* (commissioned for the Fowler Museum of Cultural History, Los Angeles), as well as the recent UCLA performance collaborations *Fahrenheit 451*, *Macbett*, and *The Iliad Project*.

Keywords: interactive, digital technology, digital media, performance, aesthetics

Introduction

Beyond browsing—beyond the "point and click" of mice, keyboards, and tablets—digital technology gives us the capability to make connections between people's actions, media, and physical and virtual spaces. These connections can surround and provide context for art or create it. They encourage engagement beyond basic navigation of traditionally hermetic "delivery" structures for video, audio, text, and other media. Here, we describe in detail a few media-rich interactive installations and performances developed at the HyperMedia Studio, a digital media research unit in the UCLA School of Theater, Film and Television. A set of "core technologies" that enable this work is then listed briefly. With both the creative work and technology as a backdrop, we propose components of a "digital aesthetic" that unify this work and provide a research focus for our own experimentation. Along the way, we discuss the possible extensions of this approach to museum exhibit design.

Selected Media Installation and Performance Work

...two, three, many Guevaras (1997) by Fabian Wagmister

An exploratory database installation commissioned by the Fowler Museum of Cultural History, *…two, three, many Guevaras* undertakes the challenge of analyzing the message and relevance of Latin American revolutionary Ernesto "Che" Guevara through the artworks he inspired. This constellation includes paintings, engravings, murals, posters, sculptures, poems, songs, and every other imaginable form of artistic representation. These works, originating in over thirty-three countries, amount to thousands of constituent media units. Multimedia relational database technology and large touch-sensitive displays are used to present a media-rich exploratory experience guided by an aleatory interpretative engine.

BEST COPY AVAILABLE

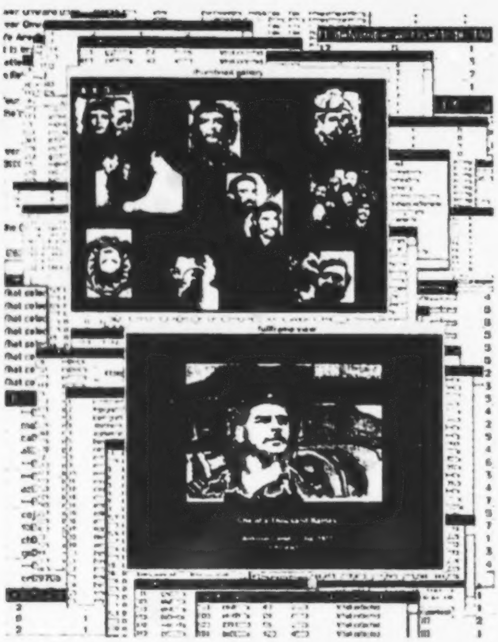

Fig. 1: Selected views from … two, three, many Guevaras

The participants are able to navigate the many creative articulations about Che and explore the complex weave of interconnections among them. The aesthetic, conceptual, historical and contextual forces informing these art works are embedded into a complex interactive navigational structure through the implementation of an adaptive search process. Rather than using standard "hyperlinks" with a single destination, the piece presents a thumbnail selection of media elements drawn from the database at random, according to a probability distribution determined by a set of interpretative modes and relevance rankings. (See Figure 1.)

These categories and weightings determine the probability that an image will appear as a thumbnail along with the one selected by the participant. This system is discussed in more detail in (Wagmister, 2000).

The participant's choices affect the weightings of each media unit, so that navigation of the piece slowly blends the author's original rankings with those generated by the participant's connection-making choices. As a result, the media selections and combinations have been significantly different at the piece's different showings; for example, in Los Angeles and Cuba. By involving the participants in a part of the creative process, the piece reflects back the viewing context in its own navigational structure.

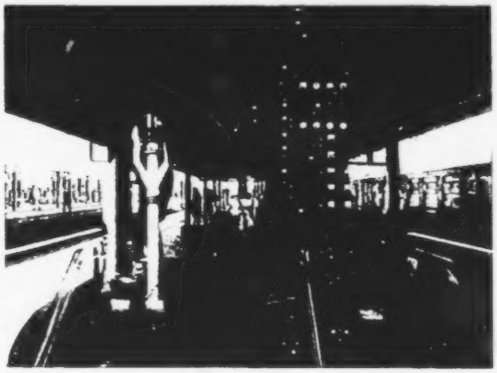

***Fig. 2: Sample image from Time&Time Again…** showing museum visitors' silhouettes filled with images of technology and superimposed on remote video. This display is output to a large rear projection screen in the museum space.*

Time&Time Again… (1999)
by Lynn Hershman and Fabian Wagmister

Time&Time Again… extends media navigation to a site-specific context with both Web- and body-based interfaces. A distributed interactive installation, the piece was commissioned by the Wilhelm Lehmbruck Museum in Duisburg, Germany, as part of its *Connected Cities* exhibit. It explores the complex relationships between our increasingly interlinked bodies and machines, and the resulting techno-cultural identity.

The installation places museum visitors and Internet viewers in a complex web of engineered interdependencies with each other and with the facilitating apparatus. At the museum, a large screen presents participants with live video images originating in train stations, coal mines, and steel mills. These locations are the nodes in a network of industrial connectivity in the Ruhr region of Germany, an infrastructure being challenged and superseded by the new network, as the industries that once defined the region become less and less cost-effective. As participants examine these projections, transfigured silhouettes of their own bodies are superimposed onto the external video signals, resulting in a composite image integrating the remote and museum elements. As each silhouette mirrors the movements of the corresponding participant, it does not reflect the details of the individual's body but instead functions as a window to a hidden scheme of technology. (See Figure 2.)

By moving toward or away from the screen, participants control the perspective of technology presented within their bodies. As in *...two, three, many Guevaras*, a large database of media is used, though it has a fixed navigation scheme in which the probability of an image being chosen for a certain layer is uniformly distributed. At one end of the installation space, away from the screen, "macro" technologies such as aerial shots of transportation, communications, or industrial structures from the Ruhr region are selected; at the opposite end, closest to the screen, images show micro perspectives such as circuits, gears, and computer boards.

A robotic, telematic doll is a synthetic witness and wired voyeur in the museum space. Simultaneously a human look-alike and a machine wannabe, its child-like form implies an alien perspective. From a corner of the installation space this "main character" of the piece observes participants' interactions with the environment and speculates on the nature of the symbiosis between humans and technology. Cameras behind its eyes stream live to the Web where remote participants at once watch and become part of the feedback system of the piece. The doll facilitates Web and museum participants' awareness of the two interconnected spaces. In addition to remote viewing, visitors to the piece's Web site can select which remote camera is displayed at the museum, control telematically the pan and zoom of the robotic doll's vision, and record video fragments into a navigable database history.

All of the dynamic media elements used in the environment are produced and combined on the fly and are not automatically recorded; each moment of the process is aleatory and ephemeral. The history database allows only Web viewers to choose to record, store, and characterize segments of the doll's vision stream, collectively creating the piece's memory.

As suggested by Eco (1989), the piece is *open*: it is completed, both conceptually and experientially, by the simultaneous action of both Web and museum visitors. It is only through the museum participants' electronically transformed silhouettes and their movement that the Web users can view the deeper layers of technology. And it is only through Web participants' actions that the images from the remote camera sites are switched and made accessible to people at the museum installation.

Invocation and Interference (2000)
by Fabian Wagmister

The museum experience of *Time & Time Again...* is driven by a sensual connection between participants and media, while the reach of human communication through technology is addressed in the online experience of the piece. These themes are revisited in a different arena in *Invocation and Interference* (*INx2*), which was first shown at "Interferences: International Festival of Electronic Art" in Belfort, France. It begins with the idea that to communicate beyond bodily reach, prevailing over the limitations of time and space, remains a constant human desire. At the personal level this need gives rise to innumerable cultural practices that regularly overlap and collide, producing unexpected readings and relational interpretations. *INx2* explores this phenomenon as experienced from a car traveling in the Argentine pampas. As one travels the immense distances of this region, two modes of very intimate communication collide in public articulation. On the one hand, the traveler encounters countless small religious shrines on the side of the road. These shrines, located in the middle of nowhere, represent promises, rememberings, gratitudes, requests to powers beyond the physical. Each shrine articulates a personal vision of popular faith and a transferring of the most intimate to the most public. At the same time, on the regional radio stations announcers regularly read personal messages destined to those who live and work in the countryside away from the reach of the telephone. These messages cover a broad set of communicational priorities, from the mundane to the tragic. For an anonymous and casual traveler, the intersection of these two communicational modes represents a significant interpretative experience.

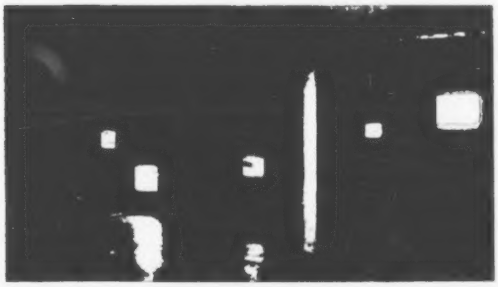

Fig. 3: The initial condition of Invocation and Interference.

Upon entering *INx2*, participants see a group of monitors of different sizes positioned on pedestals of different heights, as shown in Figure 3.

From the distance, it is clear that they reveal sections of a composite video image: the landscape of the pampas as seen from a traveling car. The sound of wind and the tuning of a car radio searching for a station is heard. Upon moving towards a particular monitor to the point where it dominates the visual field, the viewer is presented with a video view of a road shrine from a static camera position. This footage is selected at random from a large database collection. The sound of radio messages plays from that monitor at low volume, forcing the viewers to get closer to hear clearly. As the viewers move closer there is a realization that bodies control a zoom effect over the video image, allowing more careful examination of the shrine. Ultimately as the sound becomes clear, the extreme level of zoom causes the images to become distorted. Each of the monitors functions independently, allowing multiple participants to navigate the piece simultaneously. Those monitors with no person nearby continue to show the passing landscape. From a distance, the group of monitors with participants in *INx2* becomes a real time dynamic collage of the superposition of a number of forces: the driving through the pampas, the shrines of faith, the radio station's personal messages, and the interaction of the local participants.

The participants in *INx2*, by moving their bodies to explore the piece, create a sort of collective choreography with the video images for those watching from behind. David Saltz (1997) has argued that all computer-based participatory experiences are inherently *performative*. The role of the user's presence in the system of an interactive piece takes on an element of performance as the participant does what is necessary to explore it. In another area of our research, we investigate the implications of digital technology for traditional performance work and for pieces like *Hamletmachine*, that are part exhibition, part installation, and part performance.

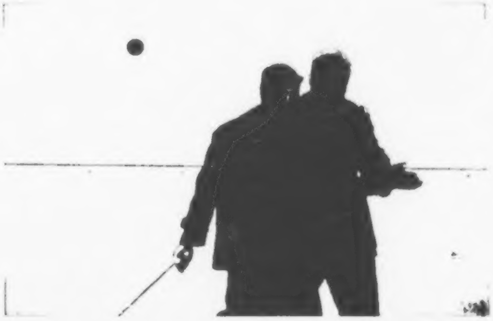

Fig. 4: A visitor to Hamletmachine experiments with the relationship between his body, shadow, and the piece's dialogue.

Hamletmachine (2000)
Installation by Jeff Burke, based on the play by Heiner Müller

In this piece, an original audio performance of German playwright Heiner Müller's *Hamletmachine* is divided into segments and stored in fifteen pieces using straightforward digital editing. In an installation space, first created for the *Fusion 2000* conference in Los Angeles, every shard is played back simultaneously, its continual loop in time unaffected by the movements of its visitors or the state of any other sound fragment. Several bright lamps at one end of the space cast their light on a long strip of sensors at waist height on the opposing wall. A visitor, who is by action or inaction part of the performance, reveals any or all of these dialogue fragments by casting a shadow on the sensors along the wall, as seen in Figure 4.

A simple relationship is constructed by a computer hidden from view: the less light on a sensor, the louder its dialogue fragment is played through the speakers in the space. The darkness of the shadow on a particular sensor controls the volume of its line at that moment without affecting the delivery itself; each sensor contains and reveals its segment of dialogue without quite allowing complete control. The different lengths of each loop ensure that the piece will almost never be the same twice.

Müller's play itself is *Hamlet*—the play, character, and the actor—ripped apart with German history and performed in pieces. This particular work further fragments it and presents the fragments simultaneously, with no one line or time privileged over another except by choice of the observer-participant. In some ways, it attempts what Jonathan Kalb describes about Robert Wilson's production of the same piece:

> The text, in other words, was simultaneously obliterated and preserved as a monument—like the images in it of Stalin, Mao, Lenin, Marx, and like Hamlet, the Hamlet Actor and his drama. *(Kalb 1998)*

The text is both sheltered and shattered by the perfect preservation and repetition possible with digital technology, while its complementary capability for dynamic manipulation of media allows each experience to be a different collage of sound and meaning. Standing close to the strip of sensors and far from the lights, one observer-performer can only reveal a few shards of dialogue at a time, but the shadows are deep and therefore the volume of each segment is quite loud. A person standing closer to the light casts a wide shadow across many sensors, revealing all fifteen fragments at once—a cacophony as if the entire play is being performed simultaneously. In between the extremes, ducking below the sensors and extending their hands into the space, or working with another person, the observer-participants can explore many other variations of the same text.

Macbett (2001)
Directed by Adam Shive, Interactive Systems by Jeff Burke

Seeing an actor familiar with *Hamletmachine* work within the space to create a unique type of performance encouraged us to continue experimenting with "interactive" technology on stage, translating the experience of developing systems for media-rich installations to performance. The recent production of Eugene Ionesco's *Macbett* at the UCLA Department of Theater was the department's first performance to incorporate "interactive systems" that allowed lighting and sound to adapt automatically to performer position and movement. *Macbett* was produced in the process typical of large shows at UCLA. It was directed and designed by graduate students, advised by faculty, managed by department staff, with undergraduate students forming the cast and crew. Like other efforts at the HyperMedia Studio, it also involved the collaboration of students from computer science and electrical engineering, who helped to develop the technical systems concurrently with the production process (Burke, 2001).

Fig. 5: In Macbett, wireless position trackers were concealed on actors and embedded in props like this witch staff.

In *Macbett*, we concentrated on the development of a toolset for defining relationships between performer action and media on stage, specifically the lighting and sound of the performance. The system worked in concert with the production's normal crew and was not designed to replace them, but instead to augment the designers' palettes with "adaptable" media components. A wireless performer tracking system was used to monitor a total of five performers and a few props used by the characters (see Figure 5).

We developed a set of software tools that allowed large-scale theatrical lighting and sound design to adapt to actor position and movement. Six computers communicating over an ethernet network performed a variety of tasks, providing graphical interfaces, controlling lighting and sound, and "interpreting" position information to make inferences about *how* performers moved instead of just *where* they were.

Though a variety of performer-driven cues were created, the most interesting were those that did not just make complex sequences more responsive to performers, but actually showed promise of affecting the process of creating theater. For example, the primary agents of the supernatural in the play—the two witches, who also appear as Lady Macbett and her Lady-in-Waiting—were each to have their own control over the stage environment through their staffs. The first conjured thunder and lighting by raising the staff quickly in the air—the quicker and stronger the thrust, the more powerful the lightning strike—while the second witch swirled her staff to create ripples of darkness, color shifts, and the sound of whirling wind proportional to the speed of her staff. These relationships were activated at the beginning of each scene where the witches appeared and lay "dormant" until the proper action was taken, allowing the actresses to conjure them up at any point. These cues required the performers to be aware of their new capabilities on stage and to work with the director to explore how they could be most effectively used.

The Iliad Project (2002)
Performance architecture by Jeff Burke, Adam Shive, Jared Stein

Macbett was a traditional play that was quite dependent on theatrical convention. From our experience with this existing text and the traditional production process arose the desire to explore the simultaneous evolution of text, design, performance technique, and technology, rather than "attaching" or "designing in" technology to another production. This change in process brings us full circle back to the concepts and technologies found in the interactive installation work discussed initially and blends them into a performance experience. *The Iliad Project* aims to develop an architecture for an original performance work that draws its themes from Homer's *Iliad* and its context from the city where it is performed. This new work is being constructed as a process that the audience intersects and influences, not simply a single, repeated performance that uses new technology. It will merge an on-line exploration of the world of the piece with a combination of interactive galleries and performance spaces. Through careful integration of a database of audience information, sensing and image capture technology, and dynamically processed media, the piece will engage the audience-participants by modifying its own text and design elements based on the groups of people who visit the Web site and attend a particular performance.

The Iliad Project's primary technological focus is audience interaction and implication through the dynamic customization of media. Where necessary, it will also incorporate dynamic control of the production environment based on performer action, as developed for *Macbett*.

Core Technologies

It is challenging to list all of the technologies that are redefining storytelling, the experience of artwork, and everyday life, though Stephen Wilson makes an amazing attempt in his new text (Wilson, 2002). Yet in our work we can find a set of "core technologies" that enable and inspire the creative projects listed above. Most of our technological interests run parallel with "ubiquitous" and "pervasive" computing research. This field includes the development of unobtrusive sensing technologies, wireless networks, and continued miniaturization of input/output interfaces. Mark Weiser's seminal 1991 article "The Computer for the 21st Century" envisioned the home and office computing systems of the future that are becoming a reality today through this work (Weiser, 1991). An excellent survey of recent ubiquitous computing research can be found in (Abowd, 2000). Information on these technologies fills many books, journals, and Web sites. They are mentioned briefly here to provide a technological background for the creative work described above.

Instrumented objects and environments

Emerging sensing technologies enables new experiences for the museum, theme park, educational space, theater, and the home. Sensors provide the entry point into the loop of communication between humans and digital systems. They can measure touch, temperature, proximity, vibration, mechanical movement, and many other physical quantities. Their intrinsic capabilities, as well as our capability to process and interpret the data they provide, will shape our environments in the future. Individual devices are becoming smaller, cheaper, and more accurate; they are starting to shed their wires and could become as ubiquitous as dust (Kahn, 1999). Rapid, accurate localization (tracking) of a large number of wireless sensors appears to be feasible in the near future (Savvides, 2001). Simultaneously, our ability to use cameras and microphones as sources of information for digital systems is continually increasing, offering possibilities for real-time control through completely remote sensors (Favaro, 2000; Wren, 1997). As these technologies mature, the need for the keyboard, mouse, and tablet will fade; this will occur most rapidly in environments emphasizing content engagement over "productivity." In our creative work, which deals with the former, we have used sensors ranging from the simple (photocells in *Time&Time Again...* and *Hamletmachine*) to the sophisticated (ultrasonic tracking in the *Macbett* positioning system).

In addition to new sensors still under development, the commercial sensing and factory automation industry manufactures a large variety of reliable, networked systems for acquisition of sensor data. We moved early on to using these types of systems for our "standard" sensing needs because of their reliability and the wide range of products available. The Tech Museum of Innovation in San Jose has used similar systems on a larger scale with much success (Ing, 1999).

Dynamic media control

In the works described above, sensing technologies are combined with digitally controlled media that include audio, video, and still images as well as lighting, motor control, and environmental effects. Digital control allows media to be the outlet for relationships connecting the physical world—as measured by sensors—with the digital world of databases, networks, and machine intelligence. The high bandwidth and large storage requirements of digital media are becoming less costly as the technology continues to develop, pushed along by the consumer entertainment market.

For now, our focus remains primarily on media that is not computer-generated: live and recorded video, audio, and still images. This focus reflects the strengths of our School of Theater, Film and Television and our desire for maximum engagement with minimum system complexity. (However, we have begun exploring 3D modeling for pre-production and sensor data visualization.) To implement flexible media delivery, we have used a number of media control technologies also found in museum systems. For example, multi-channel MPEG-2 hardware decoder cards have become our standard method of delivering many channels of high quality full frame rate and full resolution video from a single workstation. For lighting control, we have developed custom software to control industry-standard DMX lighting systems using off-the-shelf controller hardware. For sound management, we have used Cycling 74's popular Max/MSP software package running on commodity Macintosh computers. (Burke, 2001)

Databases

Digital technology's capability for "real-time" connection-making is complemented by its ability to store and query massive amounts of information. In exhibitions, performances, and media installations, the database may be used as a repository for contextual and background information. This is familiar to anyone who has designed a large database-driven Web site. However, within an artwork or interactive experience, the database also can be considered conceptually as a work's memory, leading artists and designers to develop interesting uses of databases without worrying at first about technological specifics. Databases can store historical events in a piece or exhibit (as with the Web history module in *Time&Time Again...*), remember a user's media navigation to reflect the cultural context of the viewing public (*...two, three, many Guevaras*), or link the observer-participant's experience across a variety of physical and virtual spaces (*The Iliad Project*). Coupled with natural language processing and other artificial intelligence techniques, databases provide a key component of interactive systems that move beyond simple "one-to-one" relationships between sensor inputs and media outputs.

Distributed glue

One of the unique qualities of the digital arena is the ease with which connections can be made between components, including sensors, media controllers, and databases. Because the components (or their controllers) share a common digital representation of information, they are ultimately separated only by conventions and protocols. When these can be bridged, digital technology allows artists to set up *systems of relationships* between the physical world (as it can be measured by technology), digitally controlled elements of the experience, and purely "virtual" components. Relationships might be direct (*Macbett*), adaptive (*...two, three, many Guevaras*), and/or emergent. New telecommunications networks even allow these relationships to exist almost transparently between geographically distributed components.

However, for many working with "interactive experiences" the difficulty lies in creating that initial bridge across conventions and protocols. Experimentation with connection-making is often limited by the software available and not the sensors for input or display technologies for output. Relationships between viewer-participant action and interactive work are enabled by software systems that connect or "glue together" different components of the interactive system. Our past works have used custom software developed in a variety of programming languages and authoring environments: Macromedia Director, Visual Basic, Max/MSP, and C/C++. Within the boundaries of each work, we have created flexible systems that allow experience parameters to be changed rapidly during the development and testing process. In *Macbett*, for example, a simple lighting control language was developed to allow authoring of "dynamic cues" by non-programmers. *INx2* featured a configuration system to define different relationships between participant proximity and the media elements associated with each pedestal.

Though these were fairly flexible in their specific application, each system used a slightly different approach. To facilitate future works and encourage experimentation, we are developing a control system and associated scripting language based on our experience in creating interactive works. The two are designed to provide a consistent way for non-programmers to script interactive relationships across media boundaries, allowing databases to affect stage lighting, sensors to control video playback, participant proximity to vary sound playback, and so on. We believe the approach will have applications outside of performance, in single or multi-user interactive experiences. More information on this control system will be published in (Mendelowitz, 2002).

Once a digital bridge exists across protocols and conventions, another challenge arises: How can we develop more sophisticated relationships between these interconnected elements? Initial versions of this control system will support easily expressed relationships, but the architecture will allow the addition of fuzzy controllers, machine learning and adaptation capabilities (Marti, 1999; Not, 2000), and perhaps even experimentation with autonomous actors, as already developed by Sparacino (2000) and others.

Aesthetic framework

In cinema, there is the film stock, the moving and projected image, the camera lens, and the fixed relationship between audience and screen. Theater retains the ritual of performance: someone watches, someone performs. Despite all of the arguments about what is theatrical or cinematic, the materials of the art form define a domain of parameters for artists: the "material specificity" of the medium. Though tested, pushed, and broken repeatedly by the avant garde, this domain still provides a common ground for understanding and analysis of individual works and media forms. After the choice is made to make "a film," there are unavoidable specifics of the medium, each with certain affordances: aspect ratio, grain, sprocket holes, a collection of still images moving quickly to generate the perception of motion, and so on.

Is there a digital analogy to material specificity, a framework from which to understand its uniqueness as an arena for creation? Certainly, "the digital" is, at its most basic, an abstract mathematical world. But modern engineering has generated digital devices, systems, theories, and approaches that create an arena with particular affordances which are defined by its material and virtual specifics. Our work has led us to suggest a digital aesthetic of *context*, *presence*, and *process*, enabled by digital technology's capability to define *relationships* across modal and geographic boundaries. In light of the works and technology discussed above, we develop these ideas briefly and begin to extend our discussion beyond works *in* museums to museum experiences as a whole.

Context

> [T]here is something critically useful about electronic art, even if it does not always recognize this itself. Electronic artists negotiate between the dead hand of traditional, institutionalized aesthetic discourses and the organic, emergent forms of social communication. Electronic art is an experimental laboratory, not so much for new technology as for new social relations of communication. - McKenzie Wark (1995)

In interactive works and exhibit design, providing "context" usually implies delivering a breadth of navigable content supporting a work or collection. Digital technology is often used to provide a world of contextual information by combining the storage efficiency of databases with dynamic and/or adaptive navigation schemes. An attractive feature for informational exhibits (as well as artworks) is that it has become cost-effective, at least in terms of equipment, to provide engaging, personalized access to a large amount of content.

However, digital technology encourages the exploration of a second, more individual type of context. The digital presents the opportunity to author not just contextual *content* in the traditional sense, but contextual *systems* and *interfaces* that illustrate relationships between components. The author can create a system interconnecting actual components (or their metaphorical equivalents), and that system reflects, subverts, and comments upon relationships within the piece's subject.

We have the opportunity to explore relationships with processes that, before the digital, were impossible to create because of a chasm of modality or geographic distance. Networks extend digital systems beyond the physical boundaries of the museum space, drawing in contextual information from virtual spaces (like the Web) and from different physical spaces through remote sensors. This moves beyond navigation of prebuilt contextual content to systems that place the museum, performance, or installation space squarely within the real world and its social and natural ecosystems.

For example, by the use of sensors, networks, and media control, the traffic on Santa Monica's 405 freeway might affect the flow of media through an installation space at the J. Paul Getty Museum that overlooks it. In another place, the number of people in a public square, measured against its historical average, could control the volume of a video news broadcast in an interior space. In that space, a counter ticks off the difference like a stock quote. Would CNN's coverage on September 11, 2001 exist in silence or blaring sound? How would this change if that video were also projected in the public space? If it were projected in a public space in another country? What relationships might be revealed by exploring the "memory" of such a system? By experiencing it in real time?

In *Time&TimeAgain...*, real-time context is delivered by live cameras streaming into the museum space from throughout the local region, juxtaposed with similarly "live" actions of the on-line participants. In *The Iliad Project*, which might incorporate the two scenarios in the previous paragraph, cameras and sensors in the city where the piece is performed bring live data into the space, where it is recontextualized by the event's narrative structure.

Presence

> Even the most perfect reproduction of a work of art is lacking in one element: its presence in time and space, its unique existence at the place where it happens to be. This unique existence of the work of art determined the history to which it was subject throughout the time of its existence... The presence of the original is the prerequisite to the concept of authenticity... The authenticity of a thing is the essence of all that is transmissible from its beginning, ranging from its substantive duration to its testimony to the history which it has experienced. - *Walter Benjamin* (1935)

Digital artists have often fought the perception that their media are somehow synonymous with "perfection" by creating complex or idiosyncratic systems

209

that are difficult to reproduce. In some ways, they are seeking to develop systems with the "aura" that Benjamin (1935) suggests is lacking in mechanically reproduced art. We posit that a unique type of aura is perceived when an interactive system relies on the *participant's* unique presence with the work in space and time. As someone exploring an experience, *my* body is implicated when it enables and influences the work itself. When my action, measured by sensors and affecting the environment or memory of a piece, is clearly important to that work in some way, my "performance" with the piece creates a unique sense of authenticity of experience. In stark contrast to the typical relationship between viewer and media, digital technology enables the user to influence, adapt, and explore what had previously been mechanically reproduced independent of their presence and action: the televisual and cinematic image, the media and mechanisms within a theatrical performance, the museum audio guide that plays the same audio even if the wearer is in front of the wrong work.

Moving beyond typical computer interfaces also extends today's experience of the Web. A Web site, as customized or personalized as it can become, does not (for now) care about the user's gaze or posture or presence in space. This isn't an intrinsic quality of the Web, but of the interfaces that have become its typical mode of navigation: the mouse and keyboard. Ubiquitous computing pushes us towards alternate interfaces of gesture, spoken word, position, body language, eye contact. Unlike the keyboard and the screen, these interfaces require the conversational attention of our bodies. As they change our interaction with the computer, these interfaces create difficulties for spaces and experiences built around traditional audience / art boundaries. They force us to interact with a performance, exhibit, or installation—perhaps now an art *system* instead of an object—as we would with other people, by making noise and moving around. What is exciting in the home and in educational spaces: experiential rather than observational participation; seems threatening to the traditional experience of museums, cinema, and theater.

For years, children's and science museums have had an advantage over art museums in this regard because they construct hands-on exhibits that allow visitors to touch or manipulate objects. *(Schwarzer, 2001)*

In the rush to use digital technology for unique *personal* experiences, keeping the art museum quiet, perhaps we miss new possibilities for social interaction while experiencing art.

The wearable computer and wireless Personal Digital Assistant (PDA) are a focus of attention for many museums because they maintain a "noiseless," personal experience, much like the audio guide, while providing many of the advantages of digital media. Such devices provide the annotated and enriched experiences described by Schiele (2001) and Spohrer (1999). Schwarzer (2001) points out that this technology can also be used to engage viewers in stories about the history and context of artworks. Yet no matter how dynamic the navigation between (or even inside) the stories, this does not necessarily escape a traditional observational model. The PDA ties up the hands, fixating the body on observing and manipulating its interface instead of the artwork. Without care, the aesthetic of presence may be lost in the aesthetics of point-and-click and handwriting recognition.

> [Hypermedia] allows them to choose their own paths *through* the work. But it does not cast viewers as participants within the work itself simply by virtue of employing a hypermedia interface."(Saltz, 1997)

Can we use sensing technologies in combination with wearable computers or PDAs to immediately *cast* the observer as participant in these stories or in the works themselves? Can we make their unique presence an important part of the experience?

Process

> We have, therefore, seen that (1) 'open' works, insofar as they are *in movement*, are characterized by the invitation to *make the work* together with the author and that (2) on a wider level (as a subgenus in the *species* 'work in movement') there exist works, which though organically completed, are 'open' to a continuous generation of internal relations that the addressee must uncover and select in his act of perceiving the totality of incoming stimuli.– *Umberto Eco (1989)*

This casting of the observer as an actor / co-author / participant in digitally mediated experience requires the author's willing opening up to the "risk" of what is deemed important being superseded by what is brought in by the piece's users. It forces a different role for the author and curator. In addition to creating traditional "content," the author must define relationships that connect participant's presence and context. With this comes an emphasis not on a final product or an intrinsically complete work, but on open processes and systems into which all of these components: action as discovered by sensors, media content created previously or recorded live, and navigation systems that adapt a narrative or thematic structure to a particular user.

Pieces like ...*two, three, many Guevaras* and *Time&Time Again...* illustrate that the inclusion of users in the process of adaptation requires a different kind of authoring, but not a relinquishment of the author's voice or theme. In *Macbett*, we "risked" giving direct influence over the design to the actor, breaking apart the power relationships already discarded in the poor theatre but difficult to remove in mediatized performance settings. In return, we discovered that rich media experiences are possible without requiring mechanical accuracy from the actors. In *The Iliad Project*, we take a further step, constructing the experience of the collective audience so that it depends on their individual responses and action. We cast the uninitiated in a role that is not quite performer, but more than passive audience member or extra. In installations, performances, and museum exhibits, digital technology can enable participation in a process, not just navigation of an existing, hermetic collection or exhibit.

Conclusion

To performance, installation, and exhibit design, digital technology brings the same creative challenges: context rather than isolation, presence instead of disembodiment, and process over product. Each element of the digital aesthetic suggested here is probably more difficult to implement than its alternative. But together they relate our impression of the unique nature of "the digital" that cuts across a wide range of technologies and systems. Surprisingly, we find digital technology, viewed in this light, encourages a tendency towards imperfection, unpredictability, and openness that, in return, can bring deeper audience engagement and explore new social experiences of art.

Acknowledgements

...*two, three, many Guevaras* by Fabian Wagmister in collaboration with David Kunzle, Roberto Chile, Daniel Deutsch, Fremez, and Dara Gelof. *Time&Time Again...* by Lynn Hershman and Fabian Wagmister. Team: Joel Schonbrunn, Lior Saar, Dara Gelof, Jeff Burke, Palle Henckel, Lisa Diener, Silke Albrecht. *Invocation and Interference* by Fabian Wagmister. Installation production by Dara Gelof and Jeff Burke. *Hamletmachine* installation by Jeff Burke based on text by Heiner Müller. Dialogue – Adam Shive, Meg Ferrell. Eugene Ionesco's *Macbett* directed by Adam Shive. Scenery – Maiko Nezu; Costumes – Ivan Marquez; Lighting – David Miller; Sound – David Beaudry; Interactivity – Jeff Burke and the HyperMedia Studio; Dramaturg – Sergio Costola; Stage manager – Michelle Magaldi; Production manager – Jeff Wachtel. *The Iliad Project* – architecture by Jeff Burke, Adam Shive, Jared Stein. Advisors – Fabian Wagmister, Jose Luis Valenzuela, Edit Villareal; Developers – Eitan Mendelowitz, Joseph Kim, Patricia Lee; Web developers – Caroline Ekk, Laura Hernandez Andrade, Stephan Szpak-Fleet; Production coordinator – D.J. Gugenheim.

References

Abowd, G., Mynatt, D. (2000). Charting past, present, and future research in ubiquitous computing. *ACM Trans. on Computer Human Interaction* 7, 29-58.

Benjamin, W. (1968). The Work of Art in the Age of Mechanical Reproduction. In W. Benjamin & H. Arendt (Eds.) *Illuminations*. New York: Harcourt Brace Jovanovich, 217-251.

Burke, J. (2002). Dynamic performance spaces for theatre production. *Theatre Design & Technology* 38.

Eco, U. (1989). *The Open Work*. Cambridge, Massachusetts: Harvard University Press.

Favaro, P., Jin, H., Soatto, S. (2000) Beyond point features: Integrating photometry and geometry in space and time: A semi-direct approach to structure from motion. Tech. Report CSD-200034, UCLA Department of Computer Science.

Ing., D. S. L. (1999) Innovations in a technology museum. *IEEE Micro* 19:6, 44-52.

Kahn, J. M., R. H. Katz and K. S. J. Pister. (1999) Mobile Networking for Smart Dust. *ACM/IEEE Intl. Conf. on Mobile Computing and Networking (MobiCom 99)*, Seattle, WA, August 17-19.

Marti, P., Rizzo, A. Petroni, L, Tozzi, G., Diligenti, M. (1999). Adapting the museum: a non-intrusive user modeling approach. *Proc. of the 7th Intl. Conf. on User Modeling*, Banff, Canada, 20-24 June 1999. Wien, Austria: Springer, 311-313.

Mendelowitz, E., Burke, J. (2002) A distributed control system and scripting language for 'interactivity' in live performance." To be presented at the First International Workshop on Entertainment Computing, Makuhari, Japan.

Not, E. and Zancanaro, M. (2000). The MacroNode Approach: Mediating between adaptive and dynamic hypermedia. *Proc. of Intl. Conf. on Adaptive Hypermedia and Adaptive Web-based Systesms*, Trento, Italy, 28-30 Aug. 2000. Berlin, Germany: Springer-Verlag, 167-178.

Saltz, D. Z. (1997). The Art of Interaction: Interactivity, Performativity, and Computers. *J. Aesthetics and Art Criticism* 55, 117-127.

Savvides, A., Han, C. and Srivastava, M. (2001). Dynamic fine-grained localization of ad-hoc networks of sensors. *Proc. of 7th ACM/IEEE Intl. Conf. on Mobile Computing and Networks (MobiCom'01)*. Rome: ACM. 166-179.

Schiele, B., Jebara, T., and Oliver, N. (2001) Sensory-augmented computing: Wearing the museum's guide. *IEEE Micro* 21:3, 44-52.

Schwarzer, M. (2001) Art & Gadgetry: The Future of the Museum Visit. *Museum News* July/August 2001. Accessed via http://www.cimi.org/whitesite/Handscape_Gadgets_Schwartze.htm.

Sparacino, F., Davenport, G. and Pentland, A. (2000). Media in performance: Interactive spaces for dance, theater, circus, and museum exhibits. *IBM Systems Journal* 39, 479-510.

Spohrer, J. (1999) Information in Places. *IBM Systems Journal* 38:4, 602-628.

Wagmister, F. (2000). Modular Visions: Referents, Context, and Strategies for Database Open Media Works. *AI & Society* 14, 230-242.

Wark, M. (1995). Suck on this, planet of noise! In S. Penny (Ed.), *Critical Issues in Electronic Media* Albany, New York: State University of New York Press.

Weiser, M. The Computer for the 21st Century. *Scientific American* 265, 3, 94-104.

Wilson, S. (2002) *Information Arts*. Cambridge, Massachusetts: MIT Press.

Wren, C., Azarbayejani, A., Darrell, T. and Pentland, A. (1997). Pfinder: Real-Time Tracking of the Human Body. *IEEE Trans. on Pattern Analysis and Machine Intelligence* 19, 780-785.

Repositioning the Museum

Systematically Speaking: How Do Natural History Museum Web Sites Represent Science?

Roy Hawkey, Head of Education, The Natural History Museum, London, UK

Abstract

The first natural history museum Web sites offered little more than visitor information. Then they began to include more of the nature and scope of both collections and exhibitions. Now, they incorporate sophisticated graphic design and feature active involvement by the virtual visitor, but they can also bring the museums' scientific research work to a larger and more diverse audience. Far from being principally for children, and full of dinosaurs and dioramas, major natural history museums are characterized by a high degree of fundamental scientific research. In the Eighteenth Century, museums were central to the active creation of scientific knowledge, but we now tend to associate science exclusively with laboratories. The popular image of science – test-tube and Bunsen burner – is, for several important aspects of science, inappropriate. One such area is systematics – discovering, describing, naming, classifying organisms and identifying their evolutionary relationships – a major concern of most natural history museums. Yet few Web sites explain its significance, or even make it explicit. This paper explores the representation of taxonomy, systematics and other aspects of science on selected natural history museum Web sites, using two different but related approaches.

One uses a series of categories relating to the nature of science (derived from an evaluation of exhibitions) and applies these to each website. In essence, this approach seeks to identify and, where possible, quantify evidence of representation of: 1. Science as a human endeavour – science as a social and cultural activity, a human enterprise; 2. Scientists at work – showing what scientists actually do in the process of research; 3. The status of scientific ideas – scientific ideas as theories or models, rather than as incontrovertible fact or the revelation of truth; 4. Doubt and debate – introducing scope for further questioning and reinterpretation of evidence.; 5. Opportunities for visitors to formulate their own opinions – reflecting the social construction of science. The second concentrates on science processes and practices, the methodology and operation of science: 1. Selection of research programmes – realization that science is neither certain nor neutral means that justification for research is increasingly evident, often expressed as 'biodiversity' or as benefits to humanity; 2. Collection and analysis of data – traditional and/or contemporary methods, field and laboratory techniques; 3. Evaluation of evidence and its interpretation a perception of science as unanswered questions rather than unquestioned answers; 4. Development of models, hypotheses and theories – presenting the dynamism and fluidity of science as well as an authoritative view of current understanding; 5. Publication, debate and peer review – argument and discussion as key elements of the scientific process. The paper shows that some natural history museum Web sites are now beginning to share their passion for science, especially less fashionable areas such as systematics, and that such developments coincide with changes in views about the public understanding of science and about the roles of museums.

Keywords: science; science museums; perception of science; philosophy of web sites; natural history museums; epistemology

Introduction

When it comes to attitudes to science, there is, it appears, no such person as the 'average man or woman'. A recent study concluded from extensive sampling that, in the UK, there are indeed six types, from the informed enthusiast to the 'not for me' (Wellcome Trust / OST, 2000). Yet, regardless of interest in, aptitude for or knowledge of science, many museum visitors perceive science as a series of facts and laws discovered in laboratories (no doubt by men in white coats) and expressed in difficult technical terminology. They are likely, too, to expect museums to be full of dusty objects supported by the antediluvian opinions of expert curators expressed in obscure language (Hawkey, 2001a).

Yet, in the Eighteenth Century, museums were central to the active creation of scientific knowledge. At that time museums were centres both for generating scientific understanding through research and for promulgating that understanding, through science education. Resources fundamental to both areas of activity could be drawn from their collections. Ensuing developments, however, resulted in education becoming the province of schools, while science became confined principally to laboratories. Here, experimentation could take place under controlled conditions, with the consequence that museums came to be regarded merely as places for the storage of existing – and potentially ancient – knowledge (Arnold, 1996).

So, what are natural history museums for? Somewhere to take the children on a wet afternoon? For them to marvel at the dinosaurs? Notions such as

Museums and the Web 2002: Proceedings
Archives & Museum Informatics, 2002, p. 219

BEST COPY AVAILABLE

these, together with numerous other associated assumptions, are prevalent among the general public, visitors and non-visitors alike. What misconceptions! For, more than any other type of contemporary museum, it is the natural history museums that have maintained the museums' research role. Indeed, unlike most other museums – including, ironically, museums of the physical sciences and of technology – natural history museums are characterized by a high degree of fundamental scientific research. This research activity is reflected in statements, both intentional and incidental, of the aims of natural history museums relating to using collections to make fundamental scientific discoveries about the natural world. For example, The Natural History Museum (UK) has a mission 'to maintain and develop its collections and use them to promote the discovery, understanding, responsible use and enjoyment of the natural world' (NHM, 1996).

Museums per se have such a long-established tradition – the British Museum will shortly celebrate its 250th anniversary – that it is easy to over-estimate the age of museum Web sites. Most date from the mid-1990s; many are more recent; all are continuously and rapidly evolving. The early content of natural history museum Web sites did include something of the nature and scope of both collections and exhibitions, but much was essentially visitor information – about opening times, entrance fees and bus routes. There have subsequently, of course, been vast improvements in graphic design and in enhanced functionality, the latter increasingly featuring active involvement by the virtual visitor. Latterly, other developments of great significance have begun to showcase the museums' scientific research work, conveying fundamental ideas about the life and earth sciences to a much larger and more diverse audience.

A previous paper (Hawkey, 2001b) explored natural history museum Web sites from the perspective of a science teacher, using an embryonic evaluation strategy that, not surprisingly, put great emphasis on educational issues. The present study attempts to develop further such approaches to the analysis and philosophy and practice of Web sites. To some extent it is forced to rely upon an 'expert' understanding of natural history museums and issues in science communication, especially in relation to identifying concepts and strategies that are merely implicit. It may at least provide a foundation for the development of a more sophisticated methodology.

This paper explores the representation of science on selected natural history museum Web sites. What kind of science is evident? How is it presented (in terms of epistemology rather than aesthetics)? What insights are provided into the underlying philosophy and rationale? Are there indications of the processes of investigation and enquiry or of the interpretation of evidence? What view of scientific knowledge is presented? And, in particular, what efforts are made to explain the contemporary relevance of the science of systematics?

Making Science Explicit

What, then, do natural history museum Web sites state explicitly about their science? Table 1 shows a number of useful and informative excerpts from several Web sites, in North America, Europe and Australasia. Many of these give great insight into research practice and, occasionally, philosophy. However, the ease with which such material can be located varies considerably, from those where a clear link to 'Science' or 'Research' is evident on the home page to those who provide little more than annotated lists of departmental organization. This tendency for museums to present themselves as inherently divisional or departmental, often with no more general explanations or evocations, was a greater limiting factor in the choice of museums included in this study than was, for example, language.

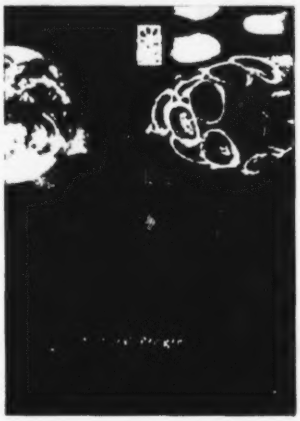

Figure 1: Natural progression

Museum	Science
American Museum of Natural History	For 125 years advanced scientific research has formed the core of the Museum. Scientists at the Museum conduct innovative research programs both in the field and within the walls of the Museum's laboratories and collections areas. (www.amnh.org)
Australian Museum	The Australian Museum undertakes an active and innovative program of research into Australia's environments and indigenous cultures. Biodiversity, geodiversity and the origins and sustainability of Australia's environments and cultures are the key focus of this research. Other work in phylogenetics concerns development and application of phylogenetic methods, philosophy of science, and editorial work for Systematic Biology. (Australian Museum, 2002)
Brussels NHM (Belgium)	...application of techniques of molecular biochemistry in the fields of systematics, phylogenesis and population genetics. (www.naturalsciences.be)
California Academy of Sciences (USA)	The California Academy of Sciences actively pursues original scientific research and is committed to fostering a spirit of scientific discovery and stewardship of the natural world. Systematic biology, the focus of the Academy's research, is becoming increasingly important as the understanding of the value of biodiversity grows. (www.calacademy.org)
Zoological Museum, Copenhagen (Denmark)	The central research areas are zoological systematics and zoogeography, including... identification, description and naming of species, interpretation of relationships (phylogeny), as well as historical and geographical aspects of evolution and biodiversity. (www.zmu.dk)
National Museum of Natural History (Smithsonian, USA)	NMNH's scientists are... filled with questions and are committed to finding the answers, have enduring curiosity, seeking and finding the puzzle pieces to significant questions about the natural world and about vital topics such as global warming, the loss of biodiversity, and invasive plant and insect species Research provides knowledge as the essential building blocks for integrative, overarching scientific interpretation. It leads to an understanding of processes that shape the natural world. The answers for today's questions come from crossing traditional academic boundaries and integrating multiple perspectives. (MNH, 2000)
The Natural History Museum (UK)	The Natural History Museum is an international leader in the scientific study of the natural world. Science is fundamental to the Museum's role [and] describes the present diversity of nature, promotes understanding of the critical importance of its past, and develops knowledge that supports anticipation and management of the impact of human activity on the environment. (www.nhm.ac.uk)
Naturalis (Netherlands)	Naturalis' collections are a source of knowledge of the characteristics and the development of the earth and life. The collections fulfill the function of natural history archive and serve as reference and basis material for research. (Naturalis, 2001a)
Royal Ontario Museum (Canada)	The ROM will be a world leader in communicating its research and collections to increase understanding of the interdependent domains of cultural and natural diversity, their relationships, significance, preservation and conservation. (www.rom.on.ca)
San Diego NHM (USA)	The extensive scientific collections... contain materials that support the research of many scientific disciplines, including those working to define and preserve biodiversity and monitor global change. (www.sdnhm.org)
Senckenberg (Frankfurt, Germany)	[holds] large collections as "documents of a nature archive". They are the fundamental basis for research activities around the world in biodiversity, in exploring the biosphere, in the evolution of life and our own origin. (www.senckenberg.de)
Swedish Museum of Natural History	Basic biological research at the Swedish Museum of Natural History concentrates on the origins of animals and plants, their systematics, and their distribution in time and space. (www.nrm.se)

Table 1: Explicit statements of science research policy and practice

A good example is provided by The Natural History Museum, where departmental designations – Botany, Entomology, Library & Information, Mineralogy, Palaeontology and Zoology – have been augmented, if not superseded, by multi-disciplinary research themes. These are presented in *Natural Progression* (NHM, 2001a), the Museum's strategy for science, which is readily available on-line as a downloadable .pdf file: figure 1.

Key elements of scientific enquiry, both process and content, are clearly discernible among the theme descriptors – as illustrated in figure 2 – and are summarized in Table 2). At the next (ie deeper) level, individual projects are highlighted in accessible language (Table 3), but with more sophisticated information and links to further work.

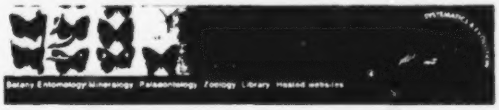
Botany Entomology Mineralogy Palaeontology Zoology Library Hosted websites

Figure 2: Making science explicit

BEST COPY AVAILABLE

Research themes	**Key aspects of science**	
	Process	content
Biomedical Sciences	Record / explain	distribution / diversity / ecology / taxonomy
Collections Management	Knowledge	
Earth Materials, History and Processes	Investigations / properties / relationships	origins / history
Ecological Patterns and Processes	Investigate / patterns	conservation / distribution
Environmental Quality	Impact assessment	aquatic & terrestrial pollution
Fauna and Flora	Description / keys / naming	diversity
Systematics and Evolution	Investigate / patterns	biodiversity / evolution / systematics

Table 2: Scientific research at The Natural History Museum (UK)

Science in the Balance

Recent developments in thinking about scientific literacy or the public understanding of science have also given increasing emphasis to the processes and practices of science (House of Lords, 2000). This remains a necessary and critical shift in emphasis, despite the fact that many formal science curricula, such as that in the UK (QCA, 1999), have begun include among their requirements some study of the nature, methodology and operation of science, in addition to an understanding of its knowledge base.

All of these issues are, to a greater or lesser extent, contained within the research programs of natural history museums. Potentially, at least, they are accessible through the material made available online. There are a variety of ways in which issues in this domain can be expressed and annotated. Table 4 shows how King's (1996) categories compare with those used by Hawkey (2001b), and indicates the synthesis that will be used in what follows in this paper.

Science as a Human Activity

How was the universe created? Are birds and dinosaurs related? What makes us human? What are the consequences of human activity on the plants and animals living in our own back yards? These and other deep questions about the natural world are what motivate our scientists each waking hour. (MNH, 2000)

The human dimension of science is critical in a number of ways, relating not only to who scientists are and what they do but also to which questions society requires them to answer. Presenting science as a social and cultural activity, as a human enterprise, may facilitate enhanced access and help to question the oft-supposed neutrality of science. That science is a human activity (rather than merely predetermined or abstract) can be reflected to some extent by repetition of 'the museum's scientists', but is best demonstrated by individualized and personalized narratives. Some Web sites, especially those of the larger US natural history museums, include scientists' names, photographs, case studies and even live links to the field. For example, Chicago's Field Museum site has on-line exhibits on 'Women in Science' (http://www.nhm.ac.uk/museum/tempexhib/gobi/gobi2.html) and on 'Adventures in the Field' (http://www.fmnh.org/exhibits/exhibit_sites/stories/map/default.htm).

Members of the general public have, in recent years increasingly come to question earlier notions that

Evolution : The Systematics and Evolution Theme	**Origin and early evolution of tetrapods (animals with four legs).** http://www.nhm.ac.uk/science/intro/palaeo/project5/index.html
The aim of the Systematics and Evolution Theme is to discover and investigate the broad patterns of biodiversity and evolution, as a foundation for comparative biology and its uses. Scientists use both traditional and modern techniques, the latter frequently derived from molecular biology, to investigate the systematics and evolution of key groups ranging from microbes to fish.	About 380 million years ago, during the Devonian period, a group of fishes evolved limbs and began to leave the water. This move was a tremendous success; all back-boned land animals, or 'tetrapods', (amphibians, reptiles, birds, and mammals, including humans) that have ever lived can ultimately trace their ancestry to these fishes.

Table 3: Example of a research theme and programme

Science communication categories (King, 1996)	Science process categories (Hawkey, 2001b)	Natural history museum Web site analysis
• science as a human endeavour	• influences on / mechanisms for selection of research programmes	• Science as a human activity
• scientists at work	• the collection and analysis of data	• Science as enquiry
• doubt and debate	• the evaluation of evidence and its interpretation	• Science as debate
• the status of scientific ideas	• the development of models, hypotheses and theories	• Science as model making
• opportunities for visitors	• publication, debate and peer review	• Science and society

Table 4: Alternative ways of thinking about nature of science issues

science is inherently beneficial and worthy of support. A heightened realization that science is neither certain nor neutral – especially in its selection of topics for research – has been a significant factor in this. Rationale for research is therefore expected to be explicit, even in apparently non-controversial areas, and natural history museum Web sites are beginning to go some way towards providing this.

The most likely to be explicitly expressed, and certainly the most frequently highlighted, is 'biodiversity'. Although a term little understood by non-specialists and absent from many school science curricula – the National Curriculum in England & Wales (QCA, 1999) has recently introduced 'sustainable development', but 'biodiversity' per se is absent – biodiversity is a theme that looms large in the realm of natural history museum Web sites. Natural history museums display a range of examples of biodiversity resources on their Web sites, with explanations that range from the elegiac to the utilitarian, from the moralistic to the homo-centric. Examples are included in table 5.

Other than biodiversity, the most common rationale given for natural history museum research is the benefits that it can offer to humanity: predicting volcanic eruptions and earthquakes, increasing food supplies, locating oil and gas reserves, maintaining and conserving natural resources. Occasionally, reference is made to economic or commercial considerations – and even, rarely, to sources of funding – but often the goal is expressed simply as that of 'better understanding'.

In ways very different from that of fictionalized drama, Web sites can also redress the stereotypical notion and show 'that science is not a list of intimidating abstractions in a textbook; it is the imaginative product of personalities who rarely conform to the stereotype of an egghead with a white coat' (Farmelo, 1992).

Australian Museum	Special emphasis has been given to the best-possible use of Museum collections in regional planning, and to the links of biodiversity assessment to sustainability and economics.
California Academy of Sciences (USA)	Petroleum geologists use collections to ascertain the identity of fossils with oil deposits; the U.S. Customs Office looks to botany collections for help in identifying imported plants; and farmers and gardeners query Academy researchers about plant pests and their biological control.
Naturalis Netherlands	About 15 years ago, more than 400 species of cichlids, a family of tropical fish, lived in Lake Victoria. Since the introduction of the Nile perch, for the benefit of the fishing industry, numerous species have disappeared forever. A unique species-rich ecosystem has become unbalanced... researchers from Naturalis accurately record how the situation has developed. (Naturalis, 2001b)
Swedish Museum of Natural History	Research on the occurrence of environmental toxins in nature and their effects on animal life is also conducted at the Museum. This is devoted to determining the geographical dispersion of toxins, as well as changes in concentrations and quantities over time.
Te Papa (New Zealand)	The key project is a major Foundation for Research Science and Technology grant-funded project to survey, describe and classify the fish fauna of the New Zealand Exclusive Economic Zone (the fourth largest in the world) and to better understand its origins and relationships. This research is discovering a new species every two to three weeks. (www.tepapa.govt.nz)

Table 5: Examples of rationales for research

BEST COPY AVAILABLE

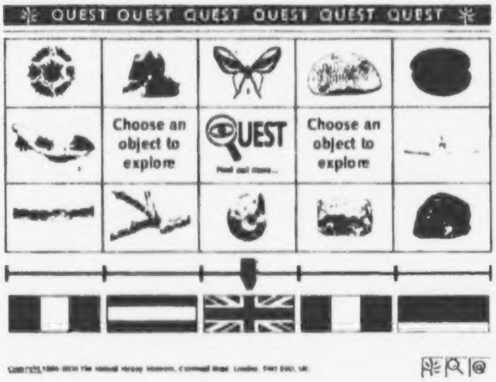
QUEST QUEST QUEST QUEST QUEST QUEST
Choose an object to explore
QUEST
Choose an object to explore

Fig. 3: QUEST's opening screen

Science as Enquiry

The vast majority of natural history museum Web sites include considerable reference to the research activity of their scientists. This may be implicit – implied by terms such as discover, describe, identify, experiment, analyse – or, more rarely, explicit. They inform visitors that scientists collect specimens from all over the world (and beyond, in the case of meteorites!) and study them using a plethora of techniques. Central to this scientific process is the collection of data and its subsequent analysis, and almost every site includes reference – whether in outline or in detail – to more traditional and/or contemporary methods. Table 6 indicates a range of these, divided into predominantly field and laboratory techniques, drawn from a variety of sites.

Some sites try to give visitors an insight in the enquiry process by involving in what is essentially a role-play activity. One such is QUEST (NHM, 1998). Here, would-be researchers are presented with a series of unidentified objects and a set of virtual tools with which they can magnify the object and look at it from different angles, find out fundamental features, such as mass, size and texture, and gather more sophisticated data including age and an image (if any) under uv light. They then make their own on-line record, with any other observations they may wish (Hawkey, 1998, 1999).

There are even examples that encourage active participation in the collection, identification and mapping of organisms such as woodlice (Hawkey, 2002a): figure 4.

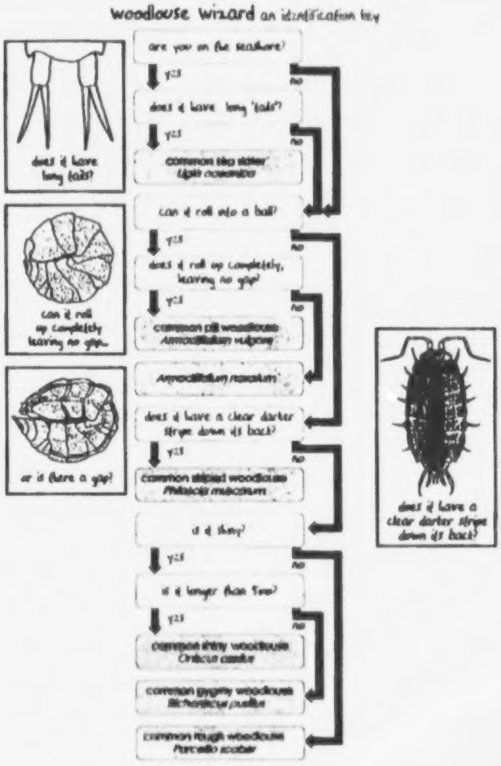
woodlouse wizard an identification key
are you on the seashore?
yes
no
does it have long 'tails'?
common sea slater Ligia oceanica
can it roll into a ball?
does it roll up completely, leaving no gap?
common pill woodlouse Armadillidium vulgare
Armadillidium nasatum
does it have a clear darker stripe down its back?
common striped woodlouse Philoscia muscorum
is it shiny?
is it longer than 5mm?
common shiny woodlouse Oniscus asellus
common pygmy woodlouse Trichoniscus pusillus
common rough woodlouse Porcellio scaber
does it have long tails?
can it roll up completely leaving no gap...
or is there a gap?
does it have a clear darker stripe down its back?

Fig. 4: An Encouragement to go 'Walking with Woodlice'

Science as Debate

Is science presented as unquestioned answers to unanswered questions (AAAS, 1993)? Do natural history museum Web sites expect their virtual visitors simply to collect knowledge or are they encouraged to engage in dialogue leading to understanding? To use a sporting analogy, does the site provide a season ticket or facilitate a free transfer (Hawkey, 2002b)?

Developing visitors' understanding, however simply, of the kinds of questions that scientists ask about evidence – and the ways in which they interpret it – should be a key aim of science communication. However, despite some clear statements of policy, the links between data collection and accepted scientific 'knowledge' are often tenuous. Presenting scientific ideas as the best model so far developed introduces scope for further questioning and rein-

BEST COPY AVAILABLE

	Laboratory	Field
Traditional	• Describe & name • Documentation • Experimental growing • Microscopic analysis	• chemical indicators • collection • ecological techniques • field observation • field research
Contemporary	• age determination by isotope analysis • computer analysis • molecular biology (DNA analysis) • 3d X-ray • mass spectrometry • high resolution transmission /scanning electron microscopy	• fossils to locate oil/gas • satellite telemetry

Table 6: Methods of data collection and analysis

terpretation of evidence. Although many sites explore the scientific research process, there is little evidence of the dynamic interplay between conflicting or competing ideas. Natural history museums as a sector have yet to acquire the confidence to expose the less committed visitor to the issues. The days when barely tolerated, ignorant visitors were expected to be grateful for whatever expert knowledge a curator was prepared to share are, somewhat paradoxically, still with us, at least on some sites. The more committed can, however, find excellent resources, often in the form of on-line essays. With its explicit emphasis on the acquisition of evidence and its subsequent, controversial interpretation, what could be more enthralling than 'Martian Meteorites, and the search for life on Mars (http://www.nhm.ac.uk/mineralogy/grady/mars.htm)' (Grady, 1999)?

There are other examples worthy of mention here. The AMNH Web site (www.amnh.org) includes many narratives about expeditions, but how telling is the admission that those to central Asia in the 1920s, which became crucial to the understanding of dinosaurs (Novacek, 1996), were actually intended to seek out the origins of humans? The new natural history museum of the Netherlands (Naturalis, 2001c) includes interesting material, too, that reflects upon the need to re-visit specimens collected 150 years ago, as species are not those attributed at the time. Owen's original description of *Iguanodon*'s thumb spike as a horn is well known, as is the discovery that *Oviraptor*'s name was inappropriately accusatory (NHM, 1997) – but at least the Web sites are not sufficiently arrogant to suggest that misinterpretation was a feature only of the naïve scientists of the past!

Science as Model Making

Clarifying the status of scientific ideas as theories or models – rather than as incontrovertible fact or the revelation of truth – can lead to a different view of scientific understanding. However, as Durant (1992) has observed of science museums in general, much of the material that is easily accessible on natural history museum Web sites (ie relatively few clicks from the home page) rather gives the impression that science is 'the sure and solid mastery of nature'. It is not that there are no reflective and discursive approaches, but that they tend to be rather deeper in the site.

Despite some attempts to indicate 'how we know' or 'what we do not yet understand', the majority of natural history museum Web sites do present science as a fixed body of knowledge. This is, in essence, little different from the perspective of their Nineteenth Century counterparts – the transmission of the curators' expert knowledge to an ignorant public. The challenge for museums is to present the dynamism and fluidity of science as well as an authoritative view of current understanding (Hawkey, 2001a). For those prepared to delve deeply, there are alternative insights. In many ways parallel to the discussion of life on Mars – and even more difficult to find – is Stringer's (1999) essay, entitled 'Were the Neanderthals Our Ancestors? (http://www.nhm.ac.uk/palaeontology/v&a/cbs/ancestors.html)' Although also concerned with evidence, this provides lively access to the nature and status of scientific ideas.

Other insights that could, given an appropriate treatment, really raise visitors' awareness of the changing status of scientific thinking are to be found on a number of sites – but often buried without further interpretation in the 'science for scientists' pages. For example, the Copenhagen Zoological Museum (2001) reports its Biosystematics Centre as having discovered and described two new animal phyla, re-evaluated hypotheses about gradients of species richness and produced a phylogenetic analysis of a new mammal species. All new ideas, overturning the old order.

Science and Society

Non-scientists – and, especially, the media – frequently express surprise when scientists disagree. And yet, whether or not one takes a Kuhnian view of paradigm shifts, argument, discussion and debate are essential components of the scientific process. Given the inherently interactive nature of the Web as a medium of communication, it is perhaps surprising that there are few examples of facilitating dialogue on museum Web sites, and certainly very few that empower visitors to formulate and contribute their own opinions. However, this is an increasing practice in the science centre sector – although not necessarily on their Web sites – and this may be expected eventually to have an impact.

There are rare examples of Web resources that allow visitors to share findings and ideas. One such is *QUEST*, significant among whose features is an online notebook. This provides access to this aspect of science – discussion and debate – that is all too often absent from conventional science resources, and certainly is rarely included as a significant component of formal education (Hawkey, 2000). Certainly, if the museums of the 21st Century are to be places for exploration and learning through discovery where – rather than provide all the answers – exhibits should be interactive and stimulate the visitor to ask questions (Abungu, 1999), then how much more should this be true of museum Web sites.

Systematically Speaking: The Science of Systematics

Science has a limited number of classical and iconic representations, paramount among which are the Bunsen burner and the test-tube, while people's awareness of science is often limited to the kind of practical, experimental work undertaken in laboratories. Although many apparently traditional laboratories may indeed be found in natural history museums, undertaking work of this type (especially in microbiology and in earth sciences), there are many other important aspects of science for which such a popular caricature is quite inappropriate. Principal among these is systematics, a major concern – indeed, the *raison d'être* – of most natural history museums: "In the progress of naming organisms and studying their relationships, systematists collect many specimens. Therefore, natural history museum collections as we know them came into being along with the science of systematics." (California Academy of Sciences, 2000).

An excellent example of a contemporary natural history museum laboratory is provided by the Pritzker (http://www.fieldmuseum.org/research_collections/pritzker_lab/pritzker/index.html) laboratory at Chicago's Field Museum, a 'core facility dedicated to genetic analysis and preservation of the world's biodiversity' (Field Museum, 2000a). The Field Museum's Pritzker Laboratory is a non-departmental multi-user core facility dedicated to genetic analysis and preservation of the world's biodiversity. It is shared by Field Museum curators, staff members and associated outside collaborators who constitute together one of the most diverse groups of evolutionary biologists and systematists in the world. The lab provides researchers with state-of-the-art equipment in molecular biology, enabling them to pursue the study of genetic diversity throughout the tree of life and at all taxonomic levels, from relationships among populations to classes and phyla of organisms.

Systematics is the science involved in the discovery, description, naming and classification of living and fossil organisms and the elucidation of their evolutionary relationships. It therefore encompasses taxonomy, the naming and classification of fossil and living species, although the two terms are often used as if they were synonyms (UK Systematics Forum, 1998). Although it constitutes a fundamental area of study in all natural history museum research, relatively few Web sites make it explicit – and even fewer explain systematics in detail or underline its significance. Those that do, however, give valuable insights into the nature of the scientific study of the natural world. As has been mentioned earlier, The Natural History Museum's Web site identifies 'sys-

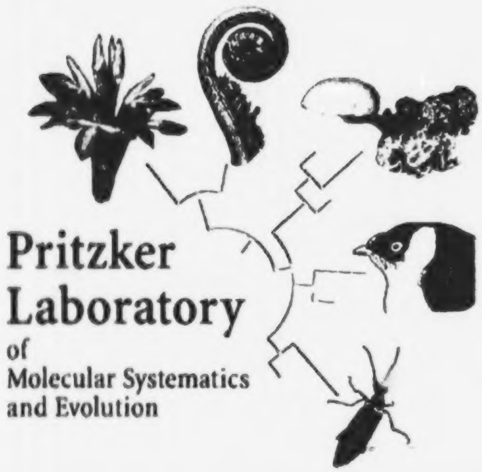

Fig. 5: An invitation to the Pritzker laboratory

tematics and evolution' as one of its key research themes.

A further component of Chicago's Field Museum of Natural History – 'Partnerships for Enhancing Expertise in Taxonomy' (http://www.fieldmuseum.org/research_collections/zoology/zoo_sites/peet/)– features not only the nature and value of the scientific work, but also the impending shortage of suitably skilled scientists (Field Museum, 2000b):

> The accelerating loss of biological diversity in the world, through habitat destruction, pollution, and ecosystem fragmentation, has been accompanied by a loss of taxonomic experts who are trained to discover, identify, describe, and classify the world's organismal diversity. Retirement of taxonomic specialists, shifts in academic recruitment and staffing, and reductions in graduate training have conjoined to impede biodiversity research and conservation, particularly on large but poorly known groups such as bacteria, fungi, protists, and numerous marine and terrestrial invertebrates. Vast numbers of species in understudied "invisible" groups constitute critical elements of food chains and ecosystems, both aquatic and terrestrial, but the high proportion of unrecognized species in these groups limits research and progress in many areas of biology and conservation.

This impending difficulty was also highlighted by an earlier UK parliamentary study (House of Lords, 1992) that bemoaned the absence of systematics from formal education courses. Recent discussions (QCA, 2002) have indicated that this omission is being addressed, although it will be several years before any changes can take effect.

Summary

Despite a number of exceptions such as those exemplified in this paper, the majority of natural history museum Web sites have yet to realize the opportunity to bring their approaches to science and science communication into the modern age. All too often science is presented only as 'revealed truth' and communication as a one-directional transmission. The potential of the Internet for museums to truly share their passion for science, especially the less fashionable areas such as systematics, is clear. That such an opportunity coincides with changes in views about the public understanding of science and about the role of museums (both already evident) makes it an opportunity not be missed. Natural history museum Web sites already provide an extensive resource, but many have some way to go before they are likely to go beyond informing the previously informed or enthusing the already enthusiastic. Most of all, they need to put less emphasis on their own internal organization and rather more on exploring the fundamental principles of science.

Taxonomy	Systematics
Taxonomy is the process of scientific description and naming of living and fossil organisms; placing them within a system of classification; and developing systems for identification. Taxonomy provides a coherent and universal system of names and is an essential foundation in any study of the natural world, particularly in the study of biological diversity and ecology.	Systematics is a broader area of scientific study. For the Museum, systematics covers the naming and investigation of the characteristics and relationships of both minerals and organisms. Systematics includes the taxonomic study of living and fossil organisms, but goes further, investigating evolution, genetics and the development of species. Systematics therefore depends both on the study of museum collections and on the investigation of variation within and between populations of organisms in the field.

Table 7: Definitions of taxonomy and systematics (UK Systematics Forum, 1998)

BEST COPY AVAILABLE

References

Abungu, G. (1999). Message from the Director-General. Nairobi: National Museums of Kenya, http://www.museums.or.ke/dg.html, visited 13 February 2002.

American Association for the Advancement of Science (1993). *Benchmarks for Scientific Literacy*, New York: Oxford University Press.

AMNH (American Museum of Natural History (2000). http://www.amnh.org/science, visited 13 February 2002.

Australian Museum (2002). www.amonline.net.au/research/index.cfm Updated January 2002, visited 13 February 2002.

Arnold, K. (1996). Presenting Science as Product or Process: Museums and the Making of Science. In Pearce, S. (Ed.), *Exploring Science in Museums*. London: Athlone Press, 57-78.

California Academy of Sciences (2000). Systematics – The Science & the Collections, http://www.calacademy.org/research/systematics.html Updated January 2002, visited 13 February 2002.

Copenhagen Zoological Museum (2001). Copenhagen Biosystematics Centre. http://www.zmuc.dk/commonweb/cobice.htm visited 13 February 2002.

Durant, J. (1992). Introduction. In Durant, J. (Ed.), *Museums and the public understanding of science*. London: Science Museum

Farmelo, G. (1992). Drama on the galleries. In Durant, J. (Ed.), *Museums and the public understanding of science*. London: Science Museum

Field Museum (2000a). Pritzker laboratory for Molecular Systematics and Evolution, http://www.fieldmuseum.org/research_collections/pritzker_lab/pritzker/index.html Updated December 2001, visited 13 February 2002.

Field Museum (2000b). Partnerships for Enhancing Expertise in Taxonomy, http://www.fieldmuseum.org/research_collections/zoology/zoo_sites/peet Updated 2000, visited 13 February 2002.

Grady, M. (1999). Martian Meteorites, and the search for life on Mars, http://www.nhm.ac.uk/mineralogy/grady/mars.htm, visited 13 February 2002.

Hawkey, R. (1998). Exploring and investigating on the Internet: virtually as good as the real thing? *Journal of Education in Museums*, 19, 16-19.

Hawkey, R. (1999). Learning from objects on-line: virtue and reality. *British Journal of Educational Technology*, 30 (1), 73-77.

Hawkey, R. (2000). Real Education from Virtual Objects: Active Learning in Science On-line. In Robson, R. (Ed.), *Proceedings of M/SET 2000, International Conference on Mathematics/Science Education & Technology*, Charlotte, VA: Association for the Advancement of Computing in Education, 202-206.

Hawkey, R. (2001a). Innovation, inspiration, interpretation: museums, science and learning. *Ways of Knowing Journal*, 1(1), 23-31.

Hawkey, R. (2001b). Science in nature and the nature of science: natural history museums on-line, *Electronic Journal of Science Education* [on-line], http://unr.edu/homepage/crowther/ejse/ejsev5n4.html, visited 13 February 2002.

Hawkey, R. (2002a). Walking with woodlice: biodiversity on-line and in the field. In Andersen, J. & Watson, D. (Eds.), *World Conference on Computers in Education VII: Networking the Learner*. Boston, MA: Kluwer Academic Publishers (in press)

Hawkey, R. (2002b). The lifelong learning game: season ticket or free transfer? *Computers and Education*, 38 (1/3), 5-20 (in press).

House of Lords (1992). *Systematic Biology Research* (First Report of the Select Committee on Science and Technology). London: HMSO

House of Lords (2000). *Science and Society* (Third Report of the Select Committee on Science and Technology). London: The Stationery Office

King, H. (1996). *The nature of science*, unpublished MSc dissertation. London: Imperial College

MNH (National Museum of Natural History) (2000). http://www.mnh.si.edu/re/index_vision.html Updated 2000, visited 13 February 2002.

Naturalis (2001a). Collections. http://www.naturalis.nl/english/collectie.html, visited 13 February 2002.

Naturalis (2001b). Eco-evolutionary biodiversity evaluation. http://www.naturalis.nl/english/biodivers.html, visited 13 February 2002.

Naturalis (2001c). Collection research. http://www.naturalis.nl/english/documentatie.html, visited 13 February 2002.

NHM (The Natural History Museum) (1996) Museum mission statement. http://www.nhm.ac.uk/info/index , visited 13 February 2002.

NHM (1997). *Dinosaurs of the Gobi desert*. http://www.nhm.ac.uk/museum/tempexhib/gobi/gobi2.html, visited 13 February 2002.

NHM (1998). *QUEST*. http://www.nhm.ac.uk/education/quest2/english/index.html , visited 13 February 2002.

NHM (2000). Science research themes, http://www.nhm.ac.uk/science/themes , visited 13 February 2002.

NHM (2001a). *Natural Progression*, London: The Natural History Museum.

NHM (2001b) Walking with woodlice. http://www.nhm.ac.uk/woodlice Updated November 2001, visited 13 February 2002.

Novacek, M. (1996). *Dinosaurs of the flaming cliffs*. New York: Anchor Books.

QCA (Qualifications and Curriculum Authority) (1999). *The National Curriculum*. London: HMSO

QCA (2002). Personal communication.

Stringer, C. (1999). Were the Neanderthals Our Ancestors? http://www.nhm.ac.uk/palaeontology/v&a/cbs/ancestors.html, visited 13 February 2002.

UK Systematics Forum (1998). *The web of life: a strategy for systematic biology in the United Kingdom*. London: UK Systematics Forum.

Wellcome Trust / Office of Science & Technology (2000). *Science and the Public: A Review of Science Communication and Public Attitudes to Science in Britain*. London: Wellcome Trust / Department of Trade & Industry.

Hacking Culture

Pia Vigh, Director CultureNet Denmark, Denmark

http://n2art.nu

Abstract

This paper argues that museums do not have a natural role in the distribution of net art, that the conservation tradition and expertise of museums do not make them suited for creating historical collections of net art without undergoing major upgrading, and that older art institutions have shown a superficial understanding of net art. Other relevant institutions already have established themselves on the Internet. Although net art does not need museums, one can still see how museums of contemporary art need net art. Public museums of contemporary art are meant to cover the whole field of contemporary art, and therefore they must necessarily also cover net art. If museums have to take net art seriously, they have to start with the already established competencies and viable forums outside the museums. Museums that wish to cover net art should join these forums. For the sake of the reputation of museums in the net art environment, it is essential that they do not appear to be parasites or "lusers" - mere users of net art who just download the resulting works of art without contributing to their structural strengthening and the more process-oriented development.

Keywords: net art; n2art; digital culture; dissemination; preservation; network art

Introduction

During the last three years, CultureNet Denmark (http://www.kulturnet.dk) along with the established cultural institutions in Denmark, has been actively working with projects digitizing and communicating cultural legacy to a broad net audience. In concert with colleagues in Sweden, Norway, Iceland and Finland, CultureNet Denmark has also created a platform for Nordic net art, n2art, hence working closely together with net artists and curators of net art. As an institution, we have learned significantly from working concurrently in the areas of institutional establishment and more anarchistic playgrounds. As individuals, we have benefited greatly from counting administrators of traditional organizations as well as unconventional avant-gardes and techno nodes among our colleagues and inspirational partners.

This paper will present the model behind the joint Nordic project of creating a platform for Nordic net art (http://n2art.nu), and continue to focus the cultural political questions of the relationship between net art and the 'art institution', which is particularly relevant to the discussion of net art – are net art and the museum incompatible operating systems? The questions revolves around a central problem at a time when archiving and communicating our digital cultural heritage is under general debate, and at a time when several works of net art have already been lost for posterity. Net art not only involves practical or technical aspects, but to a large extent also what could be called ideological aspects, which considerably complicates the discussion on communicating net art.

New media represents a constantly shifting frontier for experimentation and exploration. While new media are understood, of course, in terms of the older media that precede them, they are nonetheless freed, at least to some extent, from traditional limitations. Having to figure out how new tools work, calls for innovation and encourages a kind of beginner's mind. New media attract innovators, trendsetters and risk takers. As a result some of the hottest creative minds spend their time hacking around with new technologies that we barely understand. We need to connect and interact with these designers of the future.

Because of their newness, new media are slightly beyond the effective reach of established institutions and our bureaucracy. Net art is a case in point. While museums started to catch on to the Internet as an art medium in the last years of the nineties and began to collect, commission and exhibit net-based artworks, most artists who interest them actually

BEST COPY AVAILABLE

made their names *outside* the gallery-museum matrix. This reluctance of establishment and the freedom of the net art come at a cost. Galleries and museums serve a very important interpretive function. Museums have the qualifying ability to focus the attention of critics and audiences, to situate the artwork in a historical context, to allocate time and space for us to experience the work itself, and – not least – to preserve and protect what will be the cultural heritage of tomorrow.

Even so we are also faced with the paradox, that the Internet is the perfect tool to bring down all hierarchy and bring art and culture directly to the audience. Neither the artist nor the art institutions or the viewer have the limited roles of the past any more. What role are we to play, then, as institutions, as artists, as audience? Art – and in a general sense culture – has always been bound up with new technology, and artists have always been among the first to adopt new technologies as they emerge. Still, it sometimes seems, that the technological frontiers of art making and communication of culture is a frontier the *institutions* of arts and culture fear to tread.

The paradox of net art

In this paper, net art is understood as art, which is made for, by use of, on and/or with the Internet as the decisive technical or conceptual prerequisite. Net art is a largely anarchic artistic genre, which defines itself in opposition to the hierarchical art institution, which apparently 'dictates' what art is. Net art can be expressed in many different ways, and is broadly defined as network art in which communication is the central focal point, which means it is not necessarily bound to the Internet. In this regard, the focus is on works of Internet art in which the Internet is crucial for the work of art, hence raising particular questions of dissemination and presentation.

It is paradoxical that net art – by virtue of its criticism of the art institution – will always remain a part of this same institution, because – in one way or another – it is this positioning which makes it possible for us to discuss Internet art as an art form at all. And the very first idea of creating a joint Nordic platform for net art rested on this understanding.

The challenge posed by net art to the hierarchical structure of the art institution has a precursor in avant-garde art, which attempted to eliminate the distinction between art and non-art by challenging the art institution. The elimination of this distinction can only be of importance inside the institutional framework, because the framework makes the act that is taking place visible, of which Marcel Duchamp's *Fountain* (1917) is a classic example.

n2art – a platform for Nordic net art

During 2000, the five Nordic CultureNets initiated, developed and implemented a Nordic platform for net art: (http://n2art.nu)

What are the limits of an artwork that is part of an overall discourse on the impact of the Internet on cultural and societal developments? The concept of what is art becomes difficult to establish, because it is clear that net art seldom can be viewed as fully autonomous art objects.

n2art as a process and a project was not intended to formulate a solution to this problem, but to point out some of the paradoxes involved in net art and the institutionalisation of net art; The five Nordic CultureNets each act on behalf of the five Nordic ministries of culture, hence representing a highly bureaucratic and administrative structure. On the other hand, all five CultureNets thrive on the notion and concept of network; working in a decentralised zone of national culture. In a sense the concept of CultureNet is in itself a paradox between centre and periphery of cultural establishment.

n2art is a platform for net art, an exhibition of new art forms, but it is of course also a political construction, a prototype for Nordic co-operation, an experimental funding structure for net art. And n2art is a cultural statement. n2art is the first common Nordic project within the national CultureNets. The purpose of n2art is to establish a curated exhibition venue for net art in the Nordic region. The National CultureNets operate under the Ministry of Culture in the different Nordic Countries. But above all n2art is an experiment, an "Observatory of Premonition" – to borrow an expression from the Danish philosopher Søren Kierkegaard.

Background & Scope

n2art began as a digression, as a restless idea in the hearts and minds of the directors of the Scandinavian CultureNets. The CultureNet directors from Sweden, Norway and Denmark had long shared a common dream of inventing and implementing a joint Nordic project that would not only pave the way for Nordic co-operation in general – perhaps even pave the way for a Nordic CultureNet – but also be an expression of Nordic culture on the Internet. From the beginning it was very important to us to define a project that would in a way transcend national boundaries and stereotypes of national culture, as well as make use of the fact, that all Nordic CultureNets are situated in the magic and very intense field between culture and IT, between art and new technology.

So we chose net art; artworks designed to an only accessible on the Internet – the media of no nationality. And we were very grateful when Menningarnet (CultureNet Iceland) and Kuultturisampoo (CultureNet Finland), embraced the idea and joined the project - extending it from a Scandinavian project to a full-hearted Nordic project - as was the intention. This made n2art a shared project between the CultureNets in Sweden, Finland, Iceland, Norway and Denmark.

The purpose of n2art was to create an exhibition venue where the works of art are placed in a context, in which they are presented and annotated. The works of art have been selected by a group of 5 Nordic curators. As an experiment, it is the hope that a project like n2art will make a difference also in a political sense; that is by, increasing the recognition of this form of art and the necessity for supporting it; n2art was the first publicly funded net art site in the Nordic region. But above all n2art is important as a continued commitment on an institutional level to this form of art and discourse. So it was with great satisfaction, that we made it possible for n2art to live on – beyond our commitment – and develop within NIFCA, Nordic Institute for Contemporary Art founded and supported by Nordic Council of Ministers.

Organisation

From the very beginning it was decided to establish a parallel structure of the project dividing the focus of administration and funding from the focus of content. And this was done by establishing a Directors Board and a Curators Board, the latter consisting of one handpicked curator from each of the five participating nations.

- Over all project management and responsibility was assigned to CultureNet Denmark.
- Server management and development was assigned to ArtNet Norway being a sector net of CultureNet Norway.

The main purpose of this parallel structure was to secure the quality of the content and doing so by headhunting very competent curators and granting them the means and space for focusing only on which artworks to present. Curators Board was then given the exclusive responsibility of the content of n2art. Whereas Directors Board took care of management, administration and economy. This structure proved very efficient and productive. The sole project management and curatorial development is now conducted by NIFCA.

As a joint project n2art was not only founded on political and administrative consensus between 5 national CultureNet directors, but is also based on the artistic agreement by the 5 specially chosen curators, one from each country. All in all n2art is the result of the strong desire to co-operate - not only among the CultureNets and the directors, but also among the curators. One of the pioneers of net art, Roy Ascott, has rightly emphasized that 'networking' has been a symbol of the culture during the last decade of the 20th Century, and will continue to be one in the 21st Century. The project administration of n2art has the rare honour of participating in the realization of Ascott's prophesy

Impact & Experiences

Working with and managing the n2art project has offered CultureNet Denmark a unique experience in planning and conducting pan national joint projects. But mostly n2art has offered CultureNet Denmark familiarity with the field of net art, and we wished to place this experience and competence at the disposal of the Danish Ministry of Culture, the establishment of Danish Cultural institutions, the net art communities as well as the general public.

The questions faced working with n2art was mainly questions of:

- Funding structures
- Copyright
- 'Preservation'
- Dealing with net art as the cultural heritage of the future

Funding structures

In order to print a book, one needs a printing house. In comparison to the cost of establishing a printing house, the cost of printing a greater number of books is minimal. When one has established a printing house, one might as well print a thousand books. In order to operate a network server to host a work of net art of average complexity, one needs a computer, software, a good Internet connection that is always accessible, and a full-time technician. When this is established, one might as well host a thousand works of net art at an insignificantly higher cost.

If we at first wish to offer the most important net artists good opportunities, the price of offering good services to a broad selection of artists will merely be slightly higher. Therefore funding collective solutions - for example through supporting service providers, can to a great extent carry out public stimulation of net art. Hence we chose to establish a server environment for n2art, hosted by Artnet Norway.

The lesson learned is that support of net art and – artists can be conducted rather easy, cheap and effective, by providing server facilities. It can even be conducted by larger museums …

Copyright

Since net artists are not sufficiently economically supported through the consumption of their work, net artists are not that concerned with copyright. The basis for net art is an economy of *exchange*. Pecuniary economy has been introduced during the last two years but still exchange remains the most important. This exchange is interlinked with the free software community – net artists tend to use free software and net artists that develop software tend to share this with the net art community. Often source code is distributed in extensive friendship/colleague networks or even publicly for download. Software developed by the free software community is often protected by anti commercial licenses, for example by the common GNU, General Public Use license. (http://biomatic.org/text.php?id=55). Net artists dealing with this exchange are concerned about copyright only in the sense that the work stays free/part of the exchange.

'Preservation' of net art – mission impossible?

One important purpose of museums and libraries is the conservation of artefacts. While libraries conserve serially produced artefacts, art museums often work with unique artefacts and artefacts that have very few copies. In order to conserve unique artefacts, a tradition of selection is used, as well as an ideology that defends singling out some artefacts as more important than others. We cannot collect everything - if only for practical and economic reasons. Institutions that collect serially produced artefacts can afford a greater breadth and a less exclusive ideology.

The number of duplicates is seldom relevant in net art. Firstly, a duplicate can be distributed to its entire audience through the Internet. Secondly, duplicates can really be produced, and 'copies' can replace 'originals'. Technically it is possible to conserve all net art - depending on which conservation methods one chooses to use. However, the collection of digital information does not require less conservation technical expertise than the collection of artefacts. Still, this is presumably not sufficient to defend an exclusive conservation ideology similar to the one to which art museums traditionally adhere.

The tradition of conserving a selection of artefacts requires that we during our own time can see what is important in our time - and especially what will be important for the future. This is a quite arrogant idea. Perhaps future historians will focus on different tendencies in our time, depending on their own ideological framework. We know precious little about the ideology and culture of the future. Presumably, we will do our descendants a favour by trying to document our own time as broadly as possible.

As a starting point, the collection of net art should be modelled on literature rather than the visual arts. In Denmark for example, the National Library would perhaps be a more relevant collector of net art than the National Gallery. However, one has to bear in mind that the National Gallery and its relations are active in the artistic discourse in a completely different way from the National Library. Do we perhaps need the National Gallery as a guide to net art and as an institution that can pick out quality?

Net art and the museums

In a number of different ways, net art can be understood through an art institutional discourse, yet this does not mean that one can or should unproblematically and uncritically integrate net art into the institution. Net art is first and foremost a part of the context of the Internet. Therefore it is problematic when net art is exhibited in a museum, gallery, or on a museum's website, without the individual institution specifically relating to the artistic idioms of the Internet and drawing upon the discourse on the Internet/technology as part of the context of the artwork. Most often a museum's approach to net art is to view net art merely as the cultural heritage of tomorrow. Which of course is true! Yet net art cannot be fixed in time.

Although net art to a certain degree is dependent on an institutional framework, it is at the same time clear that the established art institution so far has not been capable of fulfilling this function. Works of net art cannot be institutionalised as autonomous art objects isolated from the context of the Internet, because they relate to the artistic idioms and discourse of the Internet. It is dynamic art, the significance of which arises out of the encounter between the artwork, the audience, and the context.

If one adds the institutional criticism by net art to the sluggishness of art institutions, it is not surprising that net art has developed its own institutions. In this context, one can point to private 'net art institutions' such as Rhizome in USA (http://rhizome.org/fresh/), ArtNet in Norway (http://kunst.kulturnett.no/artikkel.php?navn=artnet), and Artnode (http://www.artnode.dk/) in Denmark, as well as mailing lists such as nettime, which has grown on the Internet, as important disseminators of net art, because they support dialogue with and about net art and the Internet. Unlike the established art institution, these independent institutions are all a product of the Internet focused on net art and the impact of the Internet on cultural and societal development.

Virtual institutions such as Rhizome, ArtNet Norway, and Artnode Denmark purposefully incorporate the Internet in their cultural dissemination, because they are able to create an environment around the net art in which interested people can contribute to the discussion and exchange knowledge, and in this way participate in building and developing a network. Unlike the isolating framework of the art institution, the framework of the virtual institution can be characterized as the network that arise around net art. In other words, the Internet is not just an empty distribution and dissemination channel for net art. Instead, the Internet is of crucial importance for the art experience that is being disseminated, because the Internet is the context and reference point of both the virtual institutions and the works of art.

The task of museums is to conserve and disseminate works of art, but this task implies taking artworks out of their original context and isolating them as testimonies of times past, which is a problem for Internet art, in which the context is such an integral part of the artworks. In this regard, net art's own institutions contribute with a framework that is founded on the discourse on net art and the Internet. What is lacking then is the systematic collection and preserving for future dissemination and documentation, even if organizations such as ArtNet Norway, Artnode Denmark and Rhizome are striving at archiving the works, they present.

By focusing on the role of the established art institution in relation to net art, a project like n2art was bringing an essential discussion into the public debate. This is a discussion that not only has a practical aspect, but also an ideological aspect, which forces us to take a critical view of the art institution.

The institutions of net art

While painters and writers often are dependent on galleries and publishers to distribute their work, net art is independent of these institutions. One can distribute art on the entire Internet without recognition from a single living soul. Almost. And net artists use increasingly advanced computers to produce and distribute their work. Net art is a relatively new field, but specialized institutions for dissemination and discussion have already appeared.

Socially-based link selections

Many selections of links to net art are based on a relation to a particular social environment. Typically, the project overviews located at media labs that are operated by artists are produced by their own users. While these link selections emphasize the community of the social environment from which they originate, they also contribute to obscuring the artists who are not connected to the media centre. These link selections are not an expression of an evaluation of what is interesting and what is less interesting. Therefore one cannot criticize them from an artistic point of view, and they have a passive function in the art discourse. Example: Atelier Nord (http://anart.no/index.html)

Personal link selections

Most Internet sites, independent of their genre, feature a selection of links. Most lists of links of this kind are based on personal preferences and have been more or less haphazardly put together. These link selections often do not distinguish between what is art and what is not art, and contribute to maintaining the unity of net art with other types of Internet culture. These kinds of selections are seldom put into a critical context and remain difficult, if not impossible, to criticize. All criticism can be repudiated with reference to personal preferences. Example: John Hopkin's site (http://neoscenes.net/)

Liberally moderated link selections/guides

Unmoderated link selections are related to guides such as Yahoo and Alta Vista. They contain a large selection of links with a minimum amount of descriptive or contextualizing text. Such sites have as their purpose to map the field of net art in the way it sees itself, but do not attempt actual evaluation of quality. Such unmoderated collections contribute to describing the field and thereby establish important distinctions between 'us' and 'them'. They can be criticized on a fundamental and not particularly interesting level. Examples: VeryBusy (http://verybusy.org) and Netartmuseum (http://netartmuseum.org)

Moderated electronic discussion forums

Electronic discussion forums are comparable to critical journals, as for example the journal NU on visual arts. These forums are the most authoritative in the field of net art. Whatever they discuss and focus on is of great importance for a broad environment, and typically they have an influence on the net art environment. A text describing a work of art that is written by a critic has a greater chance of being distributed than a single announcement by the artist. Hence, we see that critics are active in the distribution as an external institution on which the artist becomes dependent. These forums are based on voluntary contributions by subscribers. We can thus hardly criticize those who operate these forums for not dealing with a particular project or field. Their focus is the sum of whatever their contributors are focusing on. The responsibility for maintaining an overview of the discourse has been destroyed. We revert to a situation based on the personal evaluation by individual contributors. Examples: Rhizome (http://rhizome.org), Nettime (http://nettime.org), Recode (http://systemx.autonomous.org/recode).

Unmoderated mailing lists and open networks

Unmoderated mailing lists and open networks are the most typical representatives of something both radically different and specific to the Internet. Through the use of such forums, a contributor can freely distribute materials to several thousand recipients. These kinds of forums are invaluable as the blood veins of net art. For someone who is not entirely familiar with the Internet and net art, this undergrowth of networks can be difficult to navigate through. It requires a great deal of time and some previous knowledge to gain full use of them. Examples: mailings lists: Rhizome_Raw and 7-11 (http://7-11.org), as well as Syndicate (http://colossus.v2.nl/syndicate/)and Nice (http://nice.x-i.nu)

Collections curated in the net art environment

Net artists have for years chosen to display a selection of links to their audiences, and the word 'curate' has been used about this process. This curating has played an important role in the development of the self awareness of the net art environment. Curators are obliged to make professional choices. Such selections can therefore be criticized, and the criticism cannot be repudiated on the grounds that the selection is based on a social environment or personal preferences. Such selections are a part of the discourse, and are probably often used as reference points, but do not have greater influence on net art than the flood of information on mailing lists and open networks. Example: n2art (http://n2art.nu)

Collections curated by offline art institutions

Paradoxically, institutions lacking in competence often have the greatest authority among the general public when it comes to identifying interesting net art. The reference of a museum to net art can be based on, and be criticized for, well-established models of art. However, the problem is that the employees of the institution view net art based on these established models, while net artists often use completely different points of reference. The museum makes the framework for net art, not just practically, but also conceptually. It is not uncommon for museums to ask curators working with traditional art to put together an exhibition of net art. The curators choose traditional artists who have created a net artwork, or ask them to create a net artwork for the first time. In this way, presentations are made that may be 'innovative' in traditional art, but which are not seen as relevant on the Internet. Hence, we have a situation in which an audience is presented with net art that is not seen as representative among net artists themselves. Example: Whitney Museum of American Art (http://www.whitney.org/exhibition/2kb/internet.html)

Conclusion

In this paper, I have argued that museums do not have a natural role in the distribution of net art, that the conservation tradition and expertise of museums do not make them suited for creating historical collections of net art without undergoing major upgrading, and that older art institutions have shown a superficial understanding of net art. Relevant institutions have already established themselves on the Internet. Adaptations of net art to the traditions of museums would be inappropriately conservational - we would risk that the special experience enjoyed by net art would be lost.

Although net art does not need museums, one can still see how museums of contemporary art need net art. Public museums of contemporary art are meant to cover the whole field of contemporary art, and therefore they must necessarily also cover net art. If museums have to take net art seriously, they have to start with the already established competencies and viable forums outside the museums. Museums that wish to cover net art should join these forums. For the sake of the reputation of museums in the net art environment, it is essential that they do not appear to be parasites or lusers (a term used in hacker environments about people who merely use - and do not contribute to - the development of open programs) eg. mere users of net art who just download the resulting works of art without contributing to their structural strengthening and the more process-oriented development.

The popular understanding of new media identifies it with the use of a computer for circulation and exhibition, rather than production. If we are to understand the effects of computerization on culture as a whole, this understanding is almost certainly too limiting. It would be logical to expect cultural forms and forms of dissemination of cultural heritage on the internet to eventually adopt the conventions and experience of net art. There is no reason to privilege the computer as a machine for the exhibition and distribution of media over the computer as a tool for media production or as a media storage device. All have the same potential to change existing cultural languages. And all have the same potential to leave culture as it is. Even if the latter is highly unlikely!

Storytelling and the Web in South African Museums

Katherine J. Goodnow and Yngvar Natland, International Museums Programme, University of Bergen, Norway,
in collaboration with **Iziko Museums of Cape Town, South Africa**

Abstract

The Iziko museums in Cape Town, South Africa in collaboration with the International Museums Studies Programme at the University of Bergen, Norway, have jointly developed a Web-based concept that combines oral storytelling with new technology to connect schools in the South and North. Awaiting funding, this project will be trail tested later this year. African storytelling traditions support communal ownership of stories and involve multiple forms of expression: mime, dance, music as well as verbal narrative. The South African project considers how the internet can be utilised to support and enhance these forms of storytelling with authors in the South and North. The paper considers parameters of community involvement in widely differing socio-economic contexts. Township involvement includes extending storytelling workshops at community centres to the Web.

Focusing on technology transfer, the paper also considers the relationship between central museums developed within the apartheid system with newly established community centres on the periphery of urban centres. It reviews the need for changes to established museums in the post-apartheid period with a particular emphasis on incorporating black history and contemporary oral history into social and cultural museums, both physical and Web-based, in South Africa.

Keywords: storytelling, traditional culture, mobile van, community museum, post-colonial museum, cross-cultural programs, PDA's, wireless

The Iziko Project

In 2001 and 2002 the Iziko Museums of Cape Town and the International Museums Program in Bergen, Norway developed and trial tested a Web-based museums project which draws on the skills and narrative structures of South African traditional storytelling. This paper, in addition to describing the particular strategies that have been chosen for this township project, offers also a background to the need for new forms of representation in South African museums and new relationships between museums, townships and a more diverse audience. The project involves the use of a mobile van in collaboration with community centres and satellite museums in outer urban and rural centres. Stories by children as well as adults are collected through digital video and disseminated through handheld devices (PDA's) as well as through the Web to computers in collaborating schools, museums and community groups within South Africa and overseas.

The paper begins with a brief history of museums, representation and segregation in Cape Town – a necessary background to understand the need to create new stories and reach new audiences also through the use of new technologies.

A short history of the museum

In Cape Town, before 1962, the major museum – The South African Museum – housed both the natural history and the cultural history collections. It lies in the centre of town, in an area that was (and in many ways still remains) a white area due to economic segregation, and that includes institutions such as the National Gallery, The National Library and Provincial Government parliament buildings. These institutions were only recently opened to black and mixed audiences, as the majority of these facilities were designated for whites only under South Africa's apartheid system.

The Museum of South Africa was, as such, an exception amongst these institutions. However, the placement of the museum in the city centre ensured that the museum remained a white-only entity. This vision of the museum as a segregated space remains as one of the major challenges facing South African

BEST COPY AVAILABLE

museums. The geographical placement of the museums are also compounded by the fact that the representations of black and white history have changed little since the apartheid system was abolished in the 1990's. The reasons for this are partly economic– with major funding going to challenges within health, welfare and housing – and partly related to a lack of new curatorial staff to force through major changes. Returning to history, black and white cultural history were further segregated in 1962 when a decision was made to move the cultural history collections to a separate building while natural history including ethnography, and therefore black cultural history, remained at the initial site. This split was reinforced by a political decision to place the cultural history collections under the governance of "white own affairs" (Davidson, 2001) .

The layout of the cultural history collections were further reinforced this split been white and black history. The exhibition was increasingly organised such that the viewer would walk through the collections in a seemingly 'natural' pattern - from the Ancient civilisations of Greece and Egypt through to Europe, particularly Northern Europe, and finally to colonised Cape Town. This organisation, through default, meant that African history, with the exception of colonial history, did not play a part in the formation of modern day Cape Town. In the original museum, today's South African Museum, a predominantly natural history museum, the ethnographic collections were organised around the dual desires of classification of material cultures, tribes, types of weapons etc., and that of documenting an authentic African rural life before it became tainted by industrial developments.

Representations of black culture were, and in many ways still are, in line with European ideas of the noble but primitive tribesman. Many of the images of the natives were timeless and idyllic. Plaster casts were made of tribal people carrying out traditional activities. These casts were then often placed in agricultural settings. They suggested produce and activity rather than hardship or effort. Family groupings, often mother-child compositions portrayed the notion of social cohesion. These were images that supported the arguments of those who sought a solution for South African in racial segregation (Goodby, 2001). Left alone, the black population was a proud and self-contained nation.

Other images, and plaster-casts, were oriented towards titillating a white audience. Plaster-casts were made of hunters, reinforcing the image and stereotype of the savage. These plaster-casts were placed in a diorama which initially could be seen only through peep-holes. This experience allowed the white viewer close contact with the 'dangers' of black Africa but ensured a safe distance behind the glass and with only one way viewing. This diorama was closed down in the year 2000 to the outrage of local audiences and tour operators. Tour operators, in fact, threatened to stop bringing tourists, who in the current economic climate were vital to the museums income, unless the rest of the ethnographic section remained as it was with its offerings of 'real' tribesmen (Lohman, 2001).

The reality of life amongst tribal groups during the time of the collecting of this material was in fact largely different. Traditional tribal land and the Reserves, where much of the anthropological material was gathered, were heavily effected by integration and urban change - not the least through the movement of labour, predominantly male, to mining camps and urban centres where labour was needed. This form of social history was not portrayed in the museum context, and the current collections and exhibitions are still heavily influenced by the attempts of capturing 'authentic' black history.

For the anthropologists at the museum this need to 'capture' black cultural history before it changed was understandable and similar activities have been undertaken elsewhere. The problem here was the lack of inclusion of more contemporary social history such as the effect of mining and segregation. Like many post-colonial societies, professional and amateur anthropologists and missionaries were often the only groups in a position to gather and document these material cultures and histories in a manner that was acceptable and understandable for broader audiences. As such these collections have value. However, the stories they tell do not fully encompass the histories of the people they represent.

One problem facing museums today is that little alternative documentation has been collected. Historians can easily find documents and artefacts of officers and industrialists, but few voices or images remain from early or mid-colonisation. However, many cultures, including indigenous Australians and

African cultures, argue quite rightly, that their histories, if not material cultures, have been preserved through other manners of documentation - through for example oral cultures including ceremonies, song, dance and storytelling. These stories, and this manner of documenting history, remain problematic within traditional concepts of history and science and the need for historical certainty.

A number of situations have arisen in post-colonial societies in which the value of these stories have had to be reviewed. Land Rights claims have led to an introduction of storytelling within the legal apparatus. Museums also are increasingly including storytelling not simply as a manifestation of cultural output but as a possible source for historic fact. This desire to reconsider the role of storytelling has led to a number of changes in the museums of Cape Town. One manifestation of this is in the change of name of the Cape Townian museums to the Iziko Museums of Cape Town. Iziko means the hearth around which stories are told. Another manifestation is in the desire to consider how various forms of storytelling can be incorporated into the museum environment and be disseminated to a wider audience - such as in the storytelling and the Web project to be described here.

Centre-periphery relations

A further piece of background needs to be added before entering the specifics of the storytelling project. To understand the need to move both the collection of material and the dissemination of the stories and exhibitions, the geographical lay-out into which the project is placed needs to be understood. Again there is a need to provide some historical background.

Cape Town today is geographically effected by the apartheid system of the 20th century which was reinforced through Western concepts of urban planning. Urban planners in South Africa were particularly interested in Corbusier's ideas of the modern city and, as was the case in many colonies, European intellectual ideas often found a greater foothold there than in their home countries. This foothold was supported both by a lack of contestation as well as political might to carry through the changes. For Cape Town urban renewal in the height of apartheid meant the creation of the 'white' city centre through the demolition of inner-city suburbs and forced removal of the inhabitants. This 'planning' or demolition effected particularly inner-city suburbs that were deemed racially mixed - such as District Six. The inhabitants of these districts were moved to the outer-lying Cape flats according to their designated colour. The apartheid system included 15 different categories of colour and race of which only one was white.

The townships of Cape Town surrounding the economic centre, grew in part through this forced removal, and in part through the arrival of rural populations for work purposes. The townships are only a 20-minute drive from the centre of town however there is comparatively little traffic between the two sites. For many white South Africans, many of the townships are deemed too dangerous, and for the large majority of township inhabitants the journey to the centre of town, is too costly and therefore seldom made, unless employment is there - such as the case of many domestic helpers. Trips for purely entertainment reasons or education, such as a museum visit, would be highly unusual for many township inhabitants. It is important to note, however, that the different townships include a diversity of inhabitants, with some townships populated by emerging middle-class social groups. The most populated townships include predominantly economically disadvantaged groups and it is to these that references are being made.

One of the major questions then for the Iziko Museums is how to reach out to these audiences. One, fairly traditional way of ensuring a broader visitor segment, is through bussing of school groups to the centre of town and Iziko has a relatively active educational program including transport solutions. This fills a stop-gap need, however, more long-term solutions need to be sought.

New audiences and new museum structures

A variety of ways of serving these communities have been considered including mobile museums, collaboration with newly established community centres, and the construction of permanent satellite museums. Each type of entity raises a particular set of issues to be worked out before even considering the inclusion and use of new technologies.

For politicians and local government satellite museums are seen as continuing the role museums have traditionally had i.e. educating the populace particularly in the areas of science - an area which local teachers have little resources to do properly. The driving thought here, in the South African scenario, is that the museum experience will encourage local youth to remain in school. Within this perspective the museum is seen as serving local audiences. However, in the current economic climate in South Africa, museums are expected to be self-sustaining and local audiences are not financially situated to cover these costs. Therefore the museums must also reach out to the tourist market as well. Local science education at a primary level and promoting tourism sit together uncomfortably but is part of the mandate as seen by a particularly important set of gatekeepers and one that needs to be addressed in any proposed project. New technology in this scenario would then be divided - in part covering local education needs and in part disseminating information to visitors and tourists.

Community groups themselves are divided over the issue of the role of satellite or community-based museums. For some, the community centre or satellite museum should ensure local employment and support local industry in a variety of ways including the sale of local produce, attracting tourism, and employment within the museum/centre. This leads to a view of the foreigner as the primary audience for the museum rather than the local population with the content of the museum being a representation of local culture for the enjoyment or education of others. In this scenario new technology will primarily have the role of broadcasting local culture to others.

For other sections of the local community, and for museum staff, cultural heritage is primary. The collection of material culture is seen, by museum staff, as a continuation of core work. Cultural heritage involves both presentation of local materials to the tourist market as has been done previously at the original museum, both within limited exhibitions and through the museum shop, but it is also seen as the instillation of pride in local history and material culture amongst local audiences.

This goal is approached not only through the presentation of local cultures and social history but also through participation in the creation of the centre. This is not a new approach. Indigenous participation has been discussed by many museums particularly since the 1978 UNESCO seminar *Preserving Indigenous Cultures: A New Role for Museums.* Participation has moved from simple consultation in the preparation of exhibitions by settler curators to active training of indigenous curators, inclusion on administrative boards and clear guidelines regarding consultation at all stages of exhibition planning. New technology in this scenario should then be used not only in the gathering and dissemination of local information by traditional museum staff but also involve local communities on a variety of levels including the creation of exhibition materials.

Iziko stories

With this variety of gatekeepers (and sponsors have not been included though they have a set of criteria of their own), the project in focus, Iziko stories, chose a route that involved cultural heritage and speaking to both local and foreign audiences. Education was met through children and community groups use of new media as well as through oral histories. More on this below. Local participation was also primary both through consultation as well as 'curating' the project itself. The project was developed both for use within a satellite museum or community centre as well as within a mobile museum framework.

The major question arising out of the mobile museum framework was the extent to which continuity could be maintained. Local community groups asked, quite rightly, what a mobile museum would leave behind. Continuity after the mobile museum left became then another variable to be addressed through the use of new technology.

Networking was a further issue. Traditional museums and simple satellite museum structures were likely to continue centre-periphery structures. An argument was made that rural connectivity was as important as that of centre-periphery contact. This is a similar argument as that made for the development of community media and the use of satellites to connect rural areas to each other rather than simply disseminate information from the centre to outlying rural areas (Molnar and Meadows, 2001).

Storytelling

The audience for the project became similar local audiences, centre audiences as well as foreign audiences. To find a universal element amongst these groups and to ensure interaction and equivalent input, a decision was made that a common element to all these groups was that of stories. Storytelling had the further advantage of being a central part in the culture of the local groups participating in the project.

Oral transmission of history was well developed in Africa over centuries. In part this is due to a lack of written record-keeping as, with the exception of Egypt, Ethiopia and the Buman, Vai and Akan peoples, there were few written cultures developed on the African continent (Ukadike, 1993). Information, history and cultural mores were passed on primarily through song and spoken word through a variety of ceremonies, family gatherings, royal courts, festivals and rites of passage. A rich cultural history was developed in this manner. Social history was furthermore recorded and kept alive in this way.

The arrival of colonisers threatened the survival of this oral tradition as Frank Ukadike argues:

> Under colonial domination, new values initiated by Western ideologies were introduced into African life, and under neo-colonialism, Africa struggled - and is still struggling - to develop distinct national cultures. ...The European presence brought with it an onslaught of alien influences from industrial nations, that is, the 'western way of life', yet this and endemic natural disasters such as drought, famine, and disease never entirely destroyed the tenacity of the old order. *(Ukadike, 1993, p.22)*

Using traditional oral structures does not mean that one is essentialising local populations. It is simply offering a framework that is understood both by local and overseas audiences. Using a high level of audio-visual material, and drawing on traditional storytelling forms, also meant that it would be well received by a variety of local audiences. Keyan Tomaselli, a South African filmmaker, argues that this strategy is also used by African filmmakers:

> Africa is still comprised of people who exhibit varying cultures of orality, semi-literacy and total literacy. As such, African directors often find themselves interfacing between oral and literate worlds. ... Orality helps explain their episodic, often disjointed, lateral narratives, which break with Hollywood linear conventions of beginning, middle and end. *(Siyolwe, 2002)*

This disjointedness also worked well within the hypertext structures of interactive narratives on the Web or CD-rom. Both here and in traditional African storytelling situations, storytellers can, and are, often interrupted in the telling of the tale. Unlike conventional Western narratives, African stories can take divergent paths depending on the audiences reaction and interaction. While computer games and the Internet have promoted amongst youth, used to these forms of technologies, an acceptance of interactive narrative - one that would be less expected and accepted if it were disseminated through video or television. Like current interactive narratives, however, the paths to be taken are not unlimited, the storyteller - the griot - still decides the limitations of the possible paths. This was important also as the stories were to appeal to a wider audience and in as much as they were to carry an important message of rewriting learnt history and readdressing previously ingrained stereotypes of African people.

The storytellers approached in the project involved both elders who could pass on particular views of history and stories that are not included in the mainstream educational curriculum, as well as youth who would be encouraged to create new stories from their own experiences. Ideally the storytelling process is an exchange between children in South Africa and children overseas. In the initial development children in the Nordic countries, particularly Norway, have been targeted. The goal here was to enable children to promote and exchange cultural traditions through storytelling, encouraging both exploration of their own cultural identities formed by urbanisation and cross-cultural influences as well as traditional practices. Promoting storytelling also met the need of motivation for learning and communicating through new media.

Workshops and PHAKAMA

New media would be a tool to be used in conjunction with workshops in which children would develop skills in, and knowledge of, storytelling. The workshops facilitate processes in which children learn, through a variety of media, to develop critical and investigative skills.

The project drew on the skills of an established community group that ran storytelling workshops - PHAKAMA. PHAKAMA is an organisation working already with cross-cultural theatre and storytelling projects in Great Britain, India and South Africa. PHAKAMA has broad experience bringing together from the Black, Asian, Cape Coloured and Afrikaans communities in Cape Town with communities in Britain - however not as yet through the use of new media. PHAKAMA uses local craftsmen and women across media to assist in running the workshops so as to enable a continuity of practice after the workshop is held and the initial organisers have moved on. These local craftsmen and women are often attached to local community groups or to schools. The Iziko Stories draws on these experiences and organisation.

The storytelling workshops are organised around everyday, iconic images available across cultures - such as a shoe. These images give the children, in groups or individually, a starting point they can relate to for their stories. The children are then encouraged to create fictional characters with a relationship to the image so as to allow for a necessary distance to reality. At the same time, however, the choice of everyday items meant that the stories reflected, most often, their intimate personal experiences as they are written within their sphere of knowledge. Through stories and the objects, the children communicate who they are and what occupies their daily life. The choice of a common icon facilitates the communication process and identification of children across borders. The children are taught how to act out these stories within traditional African storytelling forms - including performance, dance, music and painting. The stories are also written in a branching mode to allow for later viewer interaction.

A mobile unit, complete with trained professionals, arrives at the site at the end of the week of the workshops. These professionals have a set of digital equipment to record the children's stories and performances. Children are shown how the recording and editing of their stories takes place. The children also receive and view stories from other parts of South Africa and from overseas through a map exhibition.

Map exhibition

An integral part of the project is a large walk-on map exhibition that gives the students a sense of place. One map is that of South Africa building national understanding, a second is a more regional map to encourage networking and local connectivity. In addition maps of collaborating countries may be included.

The chosen map is rolled out onto the floor of the community centre or school. On a variety of sites at which stories have been collected a set of pedestals are placed. These pedestals include PDAs through which the children can receive and send their stories. The PDAs are sturdy and battery-operated and include an integrated cell phone so that there are no moveable parts that can be damaged. The pedestals are collapsible and stackable so as not to take up too much room in the van itself. The children can walk through the map and the pedestals. They can press the screens and view a selection of taped stories created by other children or elders. They can interact with the stories with a limited number of optional branches. These stories are ideally subtitled (or dubbed if necessary) in the primary language that is used at the school.

In South Africa today English is compulsory while Afrikaans has become optional for students with a non-Afrikaans background. Local languages are varied with six primary languages in use throughout the country. English then becomes a possible common language for broader communication while it is important to include and encourage at least a version of the stories in local languages to build up local networking and reinforce local identity.

The children may also send SMS messages, through an SMS portal, to the narrators with comments about the stories and shared similar experiences. This active participation of the children is necessary to ensure that they understand the processes of the communications that are happening. Finally the children upload their stories into the PDAs through a simple visual guide. Thus the child is participating in creating the exhibition itself, thereby developing an understanding of how the exhibition is built and how new technology can be utilised for communication purposes.

Addressing the issue of continuity, the project involves leaving one cell-phone per school/community group for SMS-messaging contact as well as copies of at least some of the videoed stories for schools which would be able to view this material (i.e. that have electricity, a VHS player and a TV monitor). One obvious advantage of the cell-phone is that electricity is not necessarily a requirement.

WWW

A further feature of the project involves the use of a website which those schools with access to computers and the net can continue to access. This requires a stable satellite up-link present in large parts of South Africa. Mobile coverage is also good along most highways and in urban and semi-urban centres. Where satellite coverage is not sufficient, drop-points may be used to up- and download the stories. This leads to a greater time delay in feedback.

The website involves an intuitive interface and a simple basic structure and is therefore easy to use both by children under guidance as well as the general public. Through the website children and students can access from the home page a Web-version of the map exhibition in which they can change languages. They may select different countries on the world map, they can choose particular objects which have inspired stories, and they can continue to send SMS and emailed messages to other children through the SMS portal. In addition, through the Web, they can pick up a webcam exhibition showing where the van is and receive updates on the journey of the van and the issues that are predominant in that particular region. As such this is similar other webbed journey stories carried out in other regions of the world.

The secondary function of the website is to provide an easy to use interface for maintaining and updating the content of the Web site. The Web Back Office fulfils this function. It enables people with basic computer training to upload the stories that have been recorded to the Internet. They do this by typing in necessary information about the digital recording and adding links to the recording. Afterwards all the information is automatically uploaded and stored. It will then instantly become a part of the Map Exhibition and the Webcam Map. The Back Office is also used to upload translations and subtitles to the stories.

The stories are complemented, through the Back Office, with information on regional differences through which the van travels - questions such as population density, environmental issues, history etc. which may be of use for teachers. As the database grows it will contain a large amount of material relevant to understanding the storyteller's daily life. The Back Office ensures also that the search function which offers criteria such as themes, location, languages and age of storyteller, functions so that it may be of use both for museum staff but also for researchers and teachers who which to use the portal in teaching social studies. Finally, the Back Office provides an element of training of local staff.

Summary

The Iziko Stories project addresses many of the issues raised in this paper. It is concerned with networking within regions rather than only working on a centre-periphery axis. It is concerned with active participation of children and local communities. It works towards deepening cultural heritage not in a manner that essentialises local communities but one that allows for expression of interrelatedness. It addresses the issue of learning through access and utilisation of new technologies but also giving the children a sense of space and geography. Through stories it assures universal interest as well as expression of local perspectives, local desires and local needs.

Obviously it is not sufficient on its own to redress the unbalanced representations that have developed within South African museums during the colonial and apartheid periods. For this to happen a fundamental transformation of curatorial practices needs to take place. Other voices, other stories need to move into the established museum spaces as well as the virtual spaces. The Iziko Stories project does, however, offer a start that allows people with di-

238

verging backgrounds to participate in both the transformation of museums as well as transformations in communication practices through new technologies.

The project is still seeking both collaborative overseas partners and further sponsors.

References

Davidsen, Patricia (Director for Social History, Iziko Museums of South Africa) taped interview in Cape Town, November 2001.

Godby, Michael African *Contrasts* in The South African Museum's *Lines of Sight* exhibition catalogue. The South African Museum: Cape Town, 2001

Lohman, Jack, interview, Cape Town, November 2001.

Molnar, Helen and Meadows, Michael *Songlines to Satellites: Indigenous communication in Australia, the South Pacific and Canada* Pluto Press: Sydney, 2001

Siyolwe, Web Conversations with Keyan Tomaselli in *African Screen Magazine*, 2002 www.und.ac.za/und/ccms/aafriprog/interv.html

Ukadike, Frank W. *African Films: A retrospective and a vision for the future* in *Critical Arts*, vol 7: Duban, 1993

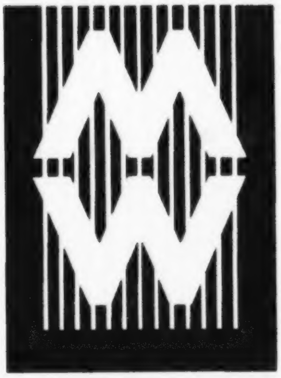

About the Authors

Authors' Biographies

Steve Allison-Bunnell has been a science writer and Web producer since the mid-1990s, when he was founding Nature and Science Editor of the Discovery Channel Online. There he produced some of the first real-world online adventures to be published on the Web, including coverage of the DCI-sponsored expeditions to the Galapagos and the Titanic. At Eduweb, he writes and produces many of science and social studies adventures. He has also written for discovery.com and britannica.com, and the U.S. Forest Service. With a B.A. in Biology from the University of Oregon and a Ph.D. in Science and Technology Studies from Cornell University, he brings a deep knowledge of the natural world to digital storytelling.

Paul M. Aoki is a Member of the Research Staff in the Computer Science Laboratory of the Xerox Palo Alto Research Center. He holds a B.S. in Electrical Engineering and an M.S. and Ph.D. in Computer Science, all from the University of California, Berkeley. His research interests include implementation techniques for database management systems as well as mobile computing and interaction design.

Nuala Bennett is the Visiting Special Projects Librarian for the IMLS-funded "Teaching with Digital Content" program, located in the Digital Imaging and Media Technology Initiative of the University of Illinois Library. She holds a B.A. in Computer Science, Linguistics and German from the University of Dublin, Trinity College (Ireland) and a M.S. in Library and Information Science from the University of Illinois. She has held appointments as Project Coordinator, "Digital Cultural Heritage Community" project, Research Information Specialist with the National Center for Supercomputing Applications (NCSA) and Research Programmer and Project Coordinator for medical informatics projects at the University of Illinois Graduate School of Library and Information Science.

Chuck Barger is Vice President of Design and Development for Interactive Knowledge, Inc., an e-learning design studio based in Charlotte, North Carolina. He has a B.A. in Philosophy from Greensboro College and an M.Ed. in Instructional Technology from the University of South Carolina. With Tim Songer, he founded Interactive Knowledge in 1991. He has been lead designer on over one hundred interactive educational titles during his career. He has presented at educational conferences across the country and is recognized as a leading expert in integrating the arts and culture into the lives of children and adults. He is currently working with national and local cultural organizations including the Smithsonian, the U.S. Department of Education as well as the Charlotte Symphony and the Mint Museum of Art.

Minda Borun is the Director of Research and Evaluation at The Franklin Institute Science Museum in Philadelphia, Pennsylvania. She has published numerous articles and four monographs on studies of visitor learning in the museum setting. She has served three terms as chair of the Committee on Audience Research and Evaluation of the American Association of Museums and was a founding member and served ten years on board of the Visitor Studies Association. Recently, Ms. Borun was lead developer for Kid Science, an NSF-funded exhibit and research project to develop a learning facility for families with children ages 5-8. Borun received her undergraduate and graduate training in Anthropology at Columbia University and has done field work among Native North Americans of the Ojibwa tribe and in a village in Mysore, South India.

Jeff Burke is a Research Fellow at the UCLA HyperMedia Studio and a Visiting Assistant Professor in the UCLA Department of Film, Television and Digital Media. He is also on the faculty of the graduate industrial design program at the Art Center College of Design in Pasadena and has undergraduate and masters degrees in Electrical Engineering from UCLA. He has designed interactive systems for numerous media installations exhibited internationally and worked in technical production in community, academic, and professional theater. His most recent collaboration with the UCLA Department of Theater was as the Interactive Systems Designer for its production of Ionesco's Macbett in the 2000-2001 subscription series.

Chris Calef graduated from the University of Oregon in 1991 with a BA in geography. He has worked as a programmer for several organizations, most recently the Oregon Center for Applied Science in Eugene, Oregon. He is skilled in database design, Internet applications and realtime 3D animation. He has traveled in Mexico, Central America, and Asia with a view toward preserving and distributing local culture and history through digital media.

Matthew Chalmers holds a BSc (Hons) Computer Science, University of Edinburgh, UK. PhD, Computer Science, U. East Anglia, UK. He serves as a Research Scientist, Rank Xerox EuroPARC, UK and as Senior Research Scientist, UBS Ubilab, Switzerland. Previously he has been Visiting Research Fellow, University of Hokkaido, Japan and is the Principal Investigator, UK EPSRC Equator Interdisciplinary Research Collaboration.

Margaret B.Chambers has an Ed.D. in Educational Media from Temple University. She is currently a consultant in educational research working with museums, zoos and other informal learning institutions. She has done extensive work on family learning in museum settings.

Timothy W. Cole is Mathematics Librarian and Associate Professor of Library Administration at the University of Illinois at Urbana-Champaign. He has been a member of the Library faculty since 1989 and is principal investigator for the Illinois Open Archives Initiative Metadata Harvesting Service project. While at Illinois he has held positions in the Beckman Institute Library, the Grainger Engineering Library, and the Library Systems Office.

Nicoletta Di Blas graduated in Classics by the Catholic University of Milan and obtained a PhD in Linguistic Sciences by the same university. She currently teaches as a research assistant Sociology of Communication by the Politecnico of Milan (campus of Como) and collaborates with the Philosophy of Language chair by the Catholic University of Milan. Her research interests focus on linguistic themes (the deictic reference in natural languages) and on the relation between rhetoric and new technologies, in particular as regards cultural heritage applications.

Andy Dingley became involved with ARKive whilst working at Hewlett Packard Labs., Bristol. This led to an interest in the Semantic Web and the other local groups involved in its development. With Paul Shabajee (Bristol University), he has published a number of papers on ARKive, with a particular emphasis on its involvement with education.

Areti Galani is a PhD student at the University of Glasgow. Holder of a Bachelor's Degree in Modern Greek Literature from the University of Athens and an MA in Museum Studies from the University of Leicester, Galani has served as Assistant Curator, Museum of the History of the University of Athens.

Katherine Goodnow is an Associate Professor in the InterMedia Department at the University of Bergen in Norway.

James M. Goodwin is a visiting professor in the UCLA Computer Science Department. He is former professor of computer science at the University of Aizu and professor emeritus of physics at California State University, Stanislaus. His Ph.D. is in physics.

Janet R. Goodwin is former professor of cultural studies at the University of Aizu, Japan. Her Ph.D. and publications are in medieval Japanese history.

Rebecca E. Grinter is a Member of the Research Staff in the Computer Science Laboratory of the Xerox Palo Alto Research Center. She holds a Ph.D. in Information and Computer Science from the University of California, Irvine. Her research interests focus on the design and evaluation of technologies to support collaboration in social and work settings.

Maria Pia Guermandi has a degree in Classical Antiquity and a Ph.D. in Archaeology from the University of Bologna Since September 1987 she has worked at the Istituto Beni Culturali of Regione Emilia Romagna where she has been involved in several research projects about applications of new technologies to cultural heritage as project leader (main areas of interest: cataloguing data bases, G.I.S., web communications). Since 1994 she has been the team leader of a research project, approved and financed by the Italian National Research Council (C.N.R.). In 2001 she joined the research team at Politecnico di Milano that analyses technology-based communication of Cultural Heritage and changing methodological paradigms, enabled or enhanced by technology. She is the author of more than 60 papers and a member of numerous Editorial Boards.

Roy Hawkey was a science teacher and university lecturer in science education before becoming Head of Education at The Natural History Museum, where he has been responsible for a number of innovations such as QUEST and Investigate. Roy has published widely on science education and given conference papers and presentations on a variety of topics, although his principal interests are the development of resources for independent and constructivist learning, the role of ICT in science education and the public understanding of science.

242

He is currently a Visiting Research Fellow in the School of Education at the University of Brighton.

Val Hobbs is senior lecturer in Information Technology at Murdoch University. Her research interests include multimedia databases, data resource management and knowledge modeling.

Jane Hunter is currently project leader of the MAENAD (Multimedia Access for Enterprises across Networks And Domains) project, which is investigating standards (Dublin Core, MPEG-7, INDECS) and schemas (RDF, XML) for multimedia resource management. She is editor of the MPEG-7 Definition Description Language (DDL) Working Draft, co-chair of the MPEG-7 DDL group, liaison between MPEG, W3Cand Dublin Core. She is also a principal investigator in the Harmony Project - an International Digital Library project between Cornell, DSTC and the University of Bristol, investigating "Metadata for the Management of Multimedia Digital Objects".

Amy Hurst is currently a third-year undergraduate at the Georgia Institute of Technology, where she is majoring in Computer Science and is a member of the Contextual Computing Group. Hurst has spent several summers as an intern at the Xerox Palo Alto Research Center. Her research interests include ubiquitous computing.

Joanne Kaczmarek is the project coordinator of the Illinois Open Archives Initiative Metadata Harvesting Service project. She received her MLS degree from the University of Illinois in May 2000. Prior to the OAI Metadata Harvesting Service project she worked for the Office of the President of the University of Illinois implementing an electronic filing system designed to digitize documents for immediate retrieval and archival purposes.

Paul F. Marty is a specialist in the application of advanced information technology to the study of history and, in particular, museum environments. He is currently the Director of Information Technology at the University of Illinois Spurlock Museum (www.spurlock.uiuc.edu) and is responsible for the design, development, and implementation of advanced information systems at this museum. He holds a BS in Computer Science Engineering, a BA in History, and an MA in Ancient History, all from the University of Illinois. He has studied Classics and Ancient History at King's College, Cambridge University, and is currently pursuing a Ph.D. in Information Science through the Graduate School of Library and Information Science at the University of Illinois.

Charles McCathieNevile is a member of the WAI domain at W3C. He is an editor of the Authoring Tool Accessibility Guidelines, and works on all four sets of WAI Guidelines. He also works on review of W3C specifications for accessibility and gives presentations and workshops on many accessibility topics, and on the Semantic Web. He has given presentations in French, Spanish and Italian as well as English. Before joining W3C in 1998 Charles worked for Sunrise Research in Melbourne on Web accessibility, Internationalization, and other projects involving educational uses of programming and the Web. He holds a degree in Medieval History from the University of Melbourne.

Slavko Milekic holds a medical degree (Belgrade School of Medicine) as well as Ph.D. in Experimental Psychology (University of Connecticut). His research interests include digital design, the psychology of human/computer interaction and building of WWW-based tools for knowledge transfer. Dr. Milekic's innovative, child-friendly digital environment (KiddyFace) has been already implemented in museum settings. He has been engaged as a consultant for the development of digital tools for art education by a number of institutions including the Philadelphia Museum of Art, Dallas Museum of Art, Please Touch Museum, Phoenix Art Museum and Cleveland Museum of Art. He currently holds the position of an associate professor of Cognitive Science & Digital Design at the University of the Arts in Philadelphia.

Libby Miller is currently a Technical Researcher at the Institute for Learning and Research Technologies (ILRT) at the University of Bristol in the UK.

Yngvar Natland is a Master Student in the Department of Media Studies at the University of Bergen in Norway

Liddy Nevile, Associate Professor, recently resigned from her position as Director of the Sunrise Research Laboratory at RMIT University in Melbourne Australia. Among other things, she has been working on projects related to the Quinkan Rock Art Interpretive centre in Laura, Queensland, and on a number of accessibility and metadata activities including working with W3C, IMS, the Education Network of Australia, the Dublin Core Metadata Initiative, the IW3C2's WWW series of conferences, and the industrial development of accessible portals for the Victorian Government in Australia.

Carolina Orsini is graduated in archaeology at the University of Bologna, and holds an MA in Cultural Heritage from the Scuola Normale Superiore in Pisa. Since 1999, she has worked in the new technologies for cultural heritage research group of the laboratory HOC and also coordinates several HOC multimedia projects. In 2000, she joined the organization of the "First Course of multimedia specialization (Internet sites, CD-rom, DVD-ROM, Virtual Worlds) for officials of museums" (organized by Institute for Cultural Heritage of the Emilia Romagna Region and from Polytechnic of Milan). She has also collaborated with different Italian museums for exhibitions and research in pre-Columbian issues.

Paolo Paolini has a degree in Physics from the University of Milan, a Masters and a Ph.D. in Computer Science from the University of California at Los Angeles (UCLA). He is presently full professor at Politecnico di Milano. He is also Lecturer at the School of Communication Sciences at the University of Italian Switzerland (USI)-Lugano. He is the scientific coordinator of HOC - Hypermedia Open Center - a team of Politecnico di Milano devoted to research and development projects concerning methods and tools for multimedia and web applications. He has more than 70 published papers, has been chairman of the ACM conference on Hypertext (1992-Milan) and co-chair of ICHIM01, and serves as Associate Editor of TOIS (Transactions on Information Systems), an ACM journal.

Darren Peacock is the Director, Information and Communication Technology, at the National Museum of Australia. Mr Peacock has worked in a variety of exhibition, marketing and IT roles in history, science and art museums. He has particular professional interests in evaluation methodologies and in the use of web technology to develop active participation in museum programs.

Diarmuid Pigott is currently a PhD student at Murdoch University, researching computerized representations of art-historical opinion. He has extensive industry experience in systems design and development, particularly in management of online digital assets.

Chris Prom is Assistant Archivist and Assistant Professor of Library Administration at the University of Illinois at Urbana-Champaign. Chris has been a member of the University of Illinois Library faculty since 1999. His research interests include the retrieval of archival metadata and archival information seeking behaviors. He also is co-principal investigator for the Illinois Open Archives Initiative Metadata Harvesting Service project.

Beth Sandore is Associate University Librarian for Information Technology Policy and Planning and Professor of Library Administration at the University of Illinois at Urbana-Champaign Library. Her professional experience and research focus on the development and evaluation of image and multimedia systems for cultural heritage information in libraries, museums and archives. She currently holds an NCSA Faculty Fellowship

Scott Sayre is the Director of Media and Technology at The Minneapolis Institute of Arts. In 1997 Scott led the formation of a collaborative online partnership with the Walker Art Center resulting in the development of ArtsConnectEd. Scott's other projects include the development of the MIA's website in 1993 and the formation of the Institute's in-house Interactive Media Group in early 1991, which has produced and installed sixteen interactive multimedia programs throughout the museum's galleries. Prior to Scott's work at the Institute, Scott held the position of Applications Researcher as the University of Minnesota's Telecommunications Development Center. Scott has a Doctorate in Education from the University of Minnesota and in 1996 received the Larry Wilson Award for Outstanding Achievement in a Non-School Environment.

David T. Schaller has more than a dozen years of experience in natural history and social science interpretation, working in print, exhibit and Web media. Before forming Educational Web Adventures in 1996, he developed exhibits for museums ranging from the Jefferson Davis Presidential Library to the Great Lakes Aquarium. At Eduweb, he has been the team leader on a wide variety of projects, ranging from colonial history adventures to space exploration simulations. He has a B.A. in Humanities from Macalester College and an M.A. in Geography and Museum Studies from the University of Minnesota.

Robert J. Semper received a Ph.D. in physics from John Hopkins University in 1973. Dr. Semper joined the Exploratorium in 1977. In 1991 he was appointed Executive Associate Director with direct responsibility for all program activities. In 1988, during a leave, he was director of a creative collaboration between Apple Computer and Lucasfilm Ltd. concerning the development of interactive multimedia education projects combining computer graphics and film and video technology. He was a Schumann Fellow and a lecturer on education in the Harvard Graduate School of Education in 1988-89. He has been a consultant for a number of museums, television programs, and educational institutions, including the Children's Television Workshop. He was the recipient of the 1994 National Science Teachers Association's Informal Educator of the Year award.

Paul Shabajee is a Research Fellow based in the Institute for Learning and Research Technology (ILRT) and the Graduate School of Education at the University of Bristol. He is working on the Hewlett Packard funded project ARKive-ERA, investigating how to design large multimedia database systems so that they are as useful as possible to a wide range of 'educational users'.

Sarah Shreeves is currently pursuing her Masters in Library and Information Science at the University of Illinois at Urbana-Champaign. She is the graduate assistant for the OAI Metadata Harvesting Project at Illinois. Before coming to Illinois, Sarah managed the successful implementation of the ALEPH 500 Library Management System at the Massachusetts Institute of Technology. Sarah holds an MA in Children's Literature from Simmons College and a BA in Medieval Studies from Bryn Mawr College.

Flavia Sparacino is a researcher at the MIT Media Lab. Her work was featured at Ars Electronica 97 and 98, ISEA 97, SIGGRAPH 96 and 99, ICMCS 99, ICHIM99, Imagina 99 and 2000, as well as many other international venues. Flavia holds undergraduate degrees in electrical engineering and robotics, masters in cognitive sciences and media arts and sciences, and a PhD from the MIT Media Lab. She received a number of scholarships and awards for her academic work and career including ones from the European Community, the Italian Center for National Research, the French Center for National Research and Fulbright, among others. Italian born, she was recently nominated Knight of the Republic of Italy, by the Italian president Carlo Azeglio Ciampi, for her contributions to art and technology.

Mirjana Spasojevic has been a researcher with Hewlett-Packard Labs since 1995 and is currently managing a research team in the CoolTown program (cooltown.hp.com), developing software architecture and infrastructure for mobile Internet services. Prior to her current position, she was a manager in the Storage Systems Program, involved with planning and configuration of enterprise storage systems. She was previously with Transarc Corporation and served as an Assistant Professor at the Washington State University School of Electrical Engineering and Computer Science. Dr. Spasojevic obtained a Ph.D. in Computer Science from The Pennsylvania State University in 1991. She shares in 4 Hewlett-Packard Co. patents and has published numerous articles and technical reports.

Margaret H. Szymanski is a Member of the Research Staff in the Systems and Practices Laboratory of the Xerox Palo Alto Research Center. She earned her Ph.D. in Spanish Linguistics from the University of California, Santa Barbara, specializing in the study of language, interaction and social organization. Her current work uses conversation analytic methods to describe how learning and knowledge sharing interactions are achieved in classrooms, museums, and workplaces.

James Thornton is a Member of the Research Staff in the Computer Science Laboratory of the Xerox Palo Alto Research Center. He holds a BMath in Computer Science from the University of Waterloo and an M.S. in Computer Science from the University of British Columbia. His research interests include distributed systems, information management, and mobile computing.

Marjolein Towler, holds an MA in Multimedia Design from Curtin University of Technology. Born in The Netherlands, Marjolein has lived in Perth, Western Australia since 1987 where she formed Consultas Pty Ltd (www.consultas.com.au) with her husband. Marjolein set up the Bachelor of Multimedia degree programme at Murdoch University from July 1998 until July 2000. She was Acting Managing Director of Imago Cooperative Multimedia Centre from September 2001 until February 2002 and Deputy Chairman from 1999 until early 2002. She is the Deputy Chair of the Film and Television Institute (FTI) in Fremantle, and National Committee Member for WA of the Australian Interactive Media Industry Association AIMIA).

Brenda Trofanenko is an Assistant Professor of Education at the University of Illinois at Urbana-Champaign. My research interests include identity and representation; public institutions and historical understanding; and relationship among academic, popular, and school history (in both digital and non-digital media). I enjoy Cole Porter songs, particularly those about wide open spaces.

Pia Vigh holds a degree in Comparative Literature from University of Copenhagen. Since 1998, she has been director of CultureNet Denmark: www.kulturnet.dk and coordinating director of the pan Nordic project: n2art - Nordic net Art. She also is a special consultant in web-dissemination.

Carl Vilbrandt is an associate professor in the Computer Arts Lab of the University of Aizu, Japan. He received his M.F.A. from the San Francisco Art Institute in San Francisco, California, USA. He has a diverse background as an art therapist, manufacturing research engineer and professional artist, sculptor and 3D computer aided designer.

Turlif Vilbrandt attended The Pittsburgh Filmmakers' School of Film, Video and Photography. He has worked as a production assistant for major Hollywood films, a videographer for the World Wildlife Fund in South America, and, for the last ten years, as a digital media consultant in his own company, Applied Vector Modeling, as well as lead programmer for the Oregon Center for Applied Science in Eugene, Oregon. He has traveled in South America, Antarctica, Alaska, Japan and Southeast Asia with a view to preserving and documenting history and nature.

Fabian Wagmister is a filmmaker and digital media artist. He is an Associate Professor in the UCLA Department of Film, Television and Digital Media and the creator of the Hypermedia Studio. Recent international exhibitions include Behind the Bars, a confrontational interactive environment about Latin America's "desaparecidos"; Time&Time Again..., with Lynn Hershman, a distributed interactive media environment exploring technological dependency and cultural identity; and the database-driven exploratory installation ...two, three, many Guevaras, examining the legacy of Ernesto Che Guevara. In these works and in his writings he combines a strong ideological voice with explorations into the protean media structures emerging from digital technologies.

Kris Wetterlund has been working with teachers as an art museum educator for the past ten years, in the education department at the Minneapolis Institute of Arts, and as Director of Education at the Minnesota Museum of American Art. Wetterlund received her degree in art education from the University of Minnesota and is certified as a K-12 Minnesota teacher. She has served as team leader in the St. Paul Public School's writing of elementary art curriculum, and has authored art educational resources online, for both the Minneapolis Institute of Arts and the Minnesota Museum of American Art. Currently Wetterlund is designing and implementing a Minnesota-wide, two year, teacher training program for ArtsConnectEd, an online partnership between The Minneapolis Institute of Arts and Walker Art Center in Minneapolis.

Michael Whittington is Curator of Pre-Columbian and African Art at the Mint Museum of Art in Charlotte, NC. He received his B.A. in anthropology from West Georgia University and an M.A. in art history from the University of Florida where he specialized in pre-Columbian art with additional coursework in African and Native American art. He conducted his graduate fieldwork in Peru and has undertaken additional research in the West Indies and southern Mexico. Michael has been a Curator at the Mint for the past five years. Among his many professional activities, Whittington co-organized the recent symposium, Praying for Rain: Style and Meaning as a Response to the Environment in Ancient American Art, at the International Congress of Americanists in Warsaw, Poland.

Allison Woodruff is a Member of the Research Staff in the Computer Science Laboratory of the Xerox Palo Alto Research Center. Woodruff holds a Ph.D. in Computer Science from the University of California, Berkeley, an M.S. in Computer Science and an M.A. in Linguistics from the University of California, Davis, and a B.A. in English from California State University, Chico. Her research interests include information visualization, user interface design, and interactive technology for historic houses.

Archives & Museum Informatics

publishes current reports from experts wordwide on issues critical to cultural heritage in the information age

Museums and the Web 2001: Selected Papers

Edited by David Bearman and
Jennifer Trant
ISBN 1-885626-23-1
230 pages, with CD-ROM of additional material.
$50.00 plus shipping and handling.

This edited paperback volume presents selected papers from **Museums and the Web 2001** together with a CD-ROM containing the full text of all other papers and demonstrations given at the conference.

Museums and the Web 2000:
Selected papers from an international conference

Edited by David Bearman and Jennifer Trant
ISBN: 1-885626-20-7
208 pages and CD-ROM
(requires Web browser)
$50.00

This print volume includes the best of the papers presented at **Museums and the Web 2000** on April 16-19th, 2000. The CD-ROM includes all papers submitted, abstracts of all presentations and biographical information for all presenters.Topics covered include Virtual Museums, Evaluation, Design Issues, and much more.

Museums and the Web 99:
Selected papers from an international conference

Edited by David Bearman and Jennifer Trant
ISBN 1-885626-17-7 (available March 1999)
245 pages, and CD-ROM (requires web browser)
$50.00

Museums and the Web 1998: Proceedings (CD-ROM)

Edited by David Bearman and Jennifer Trant
CD-ROM (requires web browser, Netscape 3.0 or equivalent) 1997. $25.00

Includes
beyond interface: net art and art on the net"
curated by Steve Dietz.

Museums and the Web, 1997: Selected Papers

Edited by David Bearman and Jennifer Trant
ISBN 1-885626-13-4 (1997), 373 pp., $30.00

International Cultural Heritage Informatics Meeting (ICHIM)

ichim

International Cultural Heritage Informatics Meeting, ichim01

edited by David Bearman, and Franca Garzotto
ISBN 1-885626-24-X
Volume 1, 655 pages
Volume 2, 519 pages .
$60.00 plus shipping and handling.

Short and full papers from the Sixth International Cultural Heritage Informatics Meeting, ichim01 in Milano, Italy, edited for this print publication. Sections of the volumes reflect the major themes discussed at ichim01: Technology, Society /Impact, Design and Evaluation, and Tools and Systems.

Cultural Heritage Informatics 1999:

Selected papers from ichim99, the International Cultural Heritage Informatics Meeting

Edited by David Bearman and Jennifer Trant
ISBN 1-885626-18-5 (1999) 255 pp., $50.00

Papers from the Fifth International Cultural Heritage Informatics Meeting, **ichim99** in Washington DC, have been edited for this print publication. Sections reflect the major themes of the conference: interactivity, converging technologies, user involvement, and new models for museum multimedia.

Museum Interactive Multimedia: Cultural Heritage Systems Design and Interfaces

Edited by David Bearman and Jennifer Trant
ISBN 1-885626-14-2 (1997), 233 pp., $30.00

A decade of progress in interactive multimedia in museums forms the basis for papers on systems design and user interface from the Fourth International Conference on Hyper-media and Interactivity in Museums (le Louvre, Paris 1997). These papers focus on design systems development and evaluation methodologies), interfaces (visitor aware systems and inter-actives providing geographical and chronological views of data), and case studies of museum ew multimedia ranging from collection catalogs to 3D environments.

Multimedia Computing and Museums

Edited by David Bearman
ISBN 1-885626-11-8 (1995) 388 pp., $20.00

Volume 1 of selected essays from the Third International Conference on Hypermedia & Interactivity in Museums (ICHIM 95 / MCN 95) on the technological, cultural and intellectual issues raised by the use of multimedia technologies to represent cultural heritage. Papers profile the iimpact of technologies on museum applications and audiences, and on the relationship of museums to society.

Hands on: Hypermedia and Interactivity in Museums

Edited by David Bearman
ISBN 1-885626-12-6 (1995) 293 pp., $20.00

Volume 2 of selected papers from the Third International Conference on Hypermedia & Interactivity in Museums (ICHIM 95 / MCN 95) reflecting the evolution of delivery mechanisms for interactive multimedia, the new social and institutional arrangements they engender, and the continuing importance of intellectual property issues. Groups of essays address fixed-format publishing, in-house interactives, networked access, museum consortia, museum teamwork, commercial partnerships and intellectual property.

Museums and Interactive Multimedia

Edited by Diane Lees
ISBN 1-8856263-89-X (1993) 436 pp., $20.00

The proceedings of the Second International Conference on Hypermedia & Interactivity in Museums (ICHIM 93) include over sixty presentations by authors from over twenty nations on the issues of design and implementation of museum interactives.

Hypermedia and Interactivity in Museums -Out of Print-

Edited by David Bearman
ISBN 1-885626-03-7 (1991) 340 pp.,

Titles on Archives and Electronic Records

Electronic Evidence: Strategies for Managing Records in Contemporary Organizations

By David Bearman

ISBN 1-885626-08-8 (1994) 314 pp., $30.00

A collection of previously published papers, accompanied by a new essay exploring the evolution of concepts of electronic records management. The papers reprinted here were originally published between 1989 and 1993, in journals in the United States, Canada, Portugal and Australia, as well as in a United Nations Report. Includes a detailed index by Victoria Irons Walch.

Order Details

Email info@archimuse.com for details or contact us at:

Archives & Museum Informatics
2008 Murray Ave, Suite D
Pittsburgh, PA 15217 USA
Phone: +1 412 422 8530
Fax: +1 412 422 8594
Email: info@archimuse.com
www.archimuse.com

Online Order Form at
http://www.archimuse.com/pub.order.html

Bulk Orders and Classroom Discounts Available

Shipping and Handling

Prepaid orders from the United States are shipped free of charge. A $5 US shipping and handling fee applies to invoiced US orders. International orders please add $10 US *per item* for shipping and handling.

Payment

All prices are in US Dollars. Payment can be made by cheque, money order, bank draft, bank transfer, or credit card (VISA, Mastercard, or American Express). Cheques must be payable to Archives & Museum Informatics, in US funds, drawn on a US Bank. Money Order or Bank Draft should be payable to Archives & Museum Informatics. Call or email for Bank Transfer details.

About the CD-ROM

This book is accompanied by a CD-ROM containing electronic versions of all these papers and many others, presented in HyperText Mark-Up Language (HTML), the technical language of the Web. It also includes abstracts of all the papers, demonstrations and workshops presented at the conference, and the biographies of all speakers and presenters. These electronic versions include color illustrations and links to the sites discussed and referenced.

You don't have to be connected to the Internet to read the papers on the CD-ROM or to navigate the full background information about the conference. You will need your own connection to the Internet to go to the linked museum sites and to follow the external links in the papers.

To use the CD-ROM, you will need a Web browser (Netscape 4.0 or Internet Explorer 4.0 or higher are recommended). Put the CD in your computer, launch your browser, and, using the File / Open menu choices, navigate to the index.html file in the main directory of the CD-ROM. Open this file in your browser – all other files are linked from there.

- **Speakers** provides a list of all the speakers at the conference and links to their abstracts, biographies and papers (where available).

- **Sessions** provides an overview of the Museums and the Web 2002 conference program and links to abstracts and paper biographies. Paper Titles link to Abstracts

- **Abstracts** describe all MW2002 program element, including demonstrations and workshops.. Titles link to full papers when available.

- **Best of the Web** will take you to the results of the Best of the Web 2002 conference, but you have to be connected to the Internet to do this.

If you have any questions or problems using the CD-ROM please email info@archimuse.com, and we'll do our best to help you.

BEST COPY AVAILABLE

Distribution of the Proceedings to
Conference Attendees sponsored by Saora Inc.

U.S. Department of Education
Office of Educational Research and Improvement (OERI)
National Library of Education (NLE)
Educational Resources Information Center (ERIC)

NOTICE

Reproduction Basis

☒ This document is covered by a signed "Reproduction Release (Blanket)" form (on file within the ERIC system), encompassing all or classes of documents from its source organization and, therefore, does not require a "Specific Document" Release form.

☐ This document is Federally-funded, or carries its own permission to reproduce, or is otherwise in the public domain and, therefore, may be reproduced by ERIC without a signed Reproduction Release form (either "Specific Document" or "Blanket").

EFF-089 (1/2003)